FORE Systems PowerHub OSPF Addendum

PN 400-1908-0001, Rev A, Issue 1, April 1996

FORE Systems reserves the right to revise this publication and to make changes in its content without obligation on the part of FORE Systems to provide notification of such revision or change. FORE Systems provides this manual without warranty of any kind, either implied or expressed, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. FORE Systems may make improvements or changes in the product(s) and/or the programs described in this manual at any time.

PowerHub is a registered trademark of FORE Systems. PowerSight and the company's logo are trademarks of FORE Systems. This document also may contain trademarks of other companies.

FORE Systems, LAN Switching Division
2115 O'Nel Drive
San Jose, CA 95131
1 408 955 9000; 1 408 955 9500 (fax)

For technical support:
email address: tech-support@alantec.com
1 800 7000 HUB; 1 408 441 0272 (fax)
Outside U.S.: contact your customer representative or call +1 408 955 9700;
+1 408 441 0272 (fax)

COMPLIANCE WITH EMISSIONS AND SAFETY STANDARDS:
This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. In addition, this product has been tested and found to be capable of operating within the specifications set forth under EN 55022 for Class A equipment. This product also complies with the provisions of EN 50082-1 relating to RFI, EMI, and ESD, which along with EN 60950 (see next paragraph) allows the product to carry the "CE" mark.

Tests performed by Inchcape Testing Services/ETL Test Laboratories demonstrate that this equipment meets UL 1950, CSA 950, and EN 60950 safety standards.

The emissions standards are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his or her own expense.

PUBLICATION HISTORY

PowerHub OSPF Addendum

PN 400-1908-0001

Date	Rev	Issue	Description
April, 1996	A	1	First edition, corresponding to the first release of OSPF routing software for the PowerHub system.

ABOUT THIS ADDENDUM

This addendum is for system administrators or others responsible for configuring the PowerHub system for OSPF routing. This addendum describes the OSPF routing protocol and the commands you use to configure your PowerHub system as an OSPF router.

This addendum contains the following chapters:

Chapter 1: Overview	Describes the OSPF routing protocol.
Chapter 2: Getting Started	Describes how to configure the PowerHub system for OSPF routing and how to troubleshoot your OSPF configuration.
Chapter 3: The OSPF Subsystem	Describes the commands in the **ospf** subsystem, which let you configure the PowerHub system to perform OSPF routing.
Chapter 4: Route Exchange Commands	Describes commands in the **rip** subsystem that let you specify how route information is exchanged between RIP and OSPF.

OTHER BOOKS

You can find additional information about your PowerHub system in the following books:

PowerHub Software Manual, V 2.6 (Rev B)
Describes the commands to configure the PowerHub 4000, 4100, 6000, and 7000 for bridging and for IP routing. This manual also describes how to configure the PowerHub system to perform RIP routing and IP Multicasting.

PowerHub Supplementary Protocols, Manual, V 2.6 (Rev B)
Describes the commands to configure the PowerHub 4000, 4100, 6000, and 7000 for AppleTalk, DECnet, and IPX routing; also describes how to implement IP Security.

PowerHub 4000 Installation and Configuration Manual, V 2.6 (Rev A)
Describes the PowerHub 4000 hardware, as well as how to install it and how to upgrade it.

PowerHub 4100 Fast Ethernet Addendum, V 41-2.6.0.1 (Rev A)
Describes the PowerHub 4100 hardware, as well as how to install it and how to upgrade it.

PowerHub 6000 Installation and Configuration Manual, V 2.6 (Rev B)
Describes the PowerHub 6000 hardware, as well as how to install it and how to upgrade it.

*PowerHub 7000 Installation and Configuration Manual, V 2.6 (*Rev C)
Describes the PowerHub 7000 hardware, as well as how to install it and how to upgrade it.

TYPOGRAPHICAL CONVENTIONS

The following typographical conventions are used in this manual:

This type style...	**Indicates...**
AaBbCcDd	A term that is being defined. Example: *OSPF (Open Shortest Path First)* is a routing protocol that enables each participating router to use a topological map of the network to route packets.
AaBbCcDd	A command name. PowerHub commands are case sensitive; they should always be entered in lower case. Example: **showcfg**
|	1) Separates the full and terse forms of a command or argument: • The full form is shown on the left of the **|**. • The terse form is shown on the right of the **|**. Example: **showcfg|scf** When you type the command or argument, you can type either the full form or the terse form. In this example, you can type **showcfg** or **scf**. 2) Separates mutually exclusive command arguments. Example: **set OSPF-router enl|dis** In this example, the command **set OSPF-router** can accept either **enl** or **dis**, but not both.
<AaBbCcDd>	Indicates a parameter for which you or the PowerHub system supplies a value. When used in command syntax, *<italics>* indicates a value you supply. Example: **set router-id** *<router-id>* In this example, *<router-id>* is a parameter for which you must supply a value when you issue this command.

`AaBbCcDd` Indicates a file name or a field name.

A file name example:

When you boot the PowerHub software, the system looks for a file named `cfg`.

A field name example:

When you boot the PowerHub software, the `login:` prompt is displayed.

`AaBbCcDd` or **`AaBbCcDd`**

Indicates text (commands) displayed by the PowerHub software or typed at the command prompt. To make typed input easy to distinguish from command prompts and output, the typed input is shown in darker type. Example:

```
1:PowerHub# pm view all 10 on 12
Port 10 (all) being viewed on: 12
```

In this example, you type "**`pm view all 10 on 12`**" and the software responds "`Port 10 (all) being viewed on: 12`".

Contents

1 Overview

This addendum describes the PowerHub **ospf** subsystem, which contains commands for configuring the PowerHub system to perform OSPF routing. In addition, this addendum describes commands in the **rip** subsystem that let you exchange route information between OSPF and RIP.

This chapter lists the PowerHub system requirements for using OSPF and describes basic features of OSPF. For complete information about OSPF, see RFC 1583. The PowerHub implementation of OSPF is based on this RFC.

- If you want to skip this overview and configure your PowerHub system as an OSPF router now, go to Chapter 2.
- If you already have configured the hub for OSPF routing but need to adjust specific OSPF parameters, go to Chapter 3.
- If you want to define filters for exchanging routes between OSPF and RIP, go to Chapter 3.

1.1 SYSTEM REQUIREMENTS

To use your PowerHub system for OSPF routing, the system must meet the following requirements.

1.1.1 Hardware

Table 1–1 lists the minimum DRAM requirements for using the PowerHub OSPF software.

TABLE 1–1 DRAM requirements for OSPF.

PowerHub Model	Packet Engine DRAM Requirement *
7000	24 MB
6000	16 MB

***** This REV of the addendum does not list the PowerHub 4000 or PowerHub 4100 DRAM requirements. DRAM requirements for the PowerHub 4000 and 4100 models will be listed in a future Rev of this addendum and in the Release Notes for those models.

To display the amount of DRAM in your Packet Engine, issue the **`mgmt showcfg`** command. (See Section 2.2 on page 12.)

The DRAM amounts listed in Table 1–1 include the DRAM requirements for all the system software, not just OSPF. You allocate a portion of the Packet Engine DRAM to the OSPF software using the **`main getmem`** command. This command allocates 500 KB of the Packet Engine DRAM to OSPF by default. However, depending upon the size of your OSPF Autonomous System, you might need less DRAM. The **`main getmem`** command lets you specify how much DRAM to allocate to OSPF. (See Section 2.5.2 on page 15.)

1.1.2 Software

To use the PowerHub system for OSPF routing your hub must have at least the software versions listed in Table 1–2.

TABLE 1–2 Software requirements for OSPF.

PowerHub Model	Minimum Software Requirements *	
	System Software	Packet Engine boot PROM
7000	7-2.6.3.0	2.5.4 or later

***** This REV of the addendum lists only the PowerHub 7000 software requirements. Software requirements for the other PowerHub models will be listed in a future Rev of this addendum and in the Release Notes for the other PowerHub models.

If you ordered your PowerHub system from the factory with one of the software versions listed in Table 1–2, you do not need to install software to use the hub for OSPF routing.

To display the version of system software running on your PowerHub system, use the **`main version`** command. (This command is described in Section 2.1 on page 11.)

If you received a version of software later than the versions listed above, install the software using the installation instructions in the Release Notes that accompany the software. After you install the software, use the procedures in Section 2.5 on page 14 to configure the hub for OSPF routing.

1.2 WHAT IS OSPF?

OSPF (Open Shortest Path First) is a routing protocol that enables each participating router to use a topological map of the network to route packets. OSPF routers exchange route information using *LSAs (link-state advertisements)*. An LSA is a packet that reports the link states (up or down) of a router's interfaces that are attached to devices in the OSPF network.

OSPF is an IGP (Interior Gateway Protocol). That is, OSPF distributes routing information within a single Autonomous System.

RIP (Routing Information Protocol) is another popular IGP used in TCP/IP networks, and also is implemented in the PowerHub software. RIP uses *distance vectors* to build route tables in RIP routers. A "hop count" (sometimes called a "cost" or a "metric") is associated with each route. The cost represents the number of hops (routers) between the packet's source host and destination host.

OSPF does not use distance vectors for routes. Instead, OSPF uses the link states of the routers in the network to derive a picture of the network topology.

OSPF lets you logically group your networks into areas. By grouping your networks into areas, you can reduce the amount of routing traffic among OSPF routers. An area's topology is summarized before being sent to other areas, thereby reducing routing traffic.

1.3 OSPF TOPOLOGY

OSPF routers view the network in terms of the following topological elements:

Autonomous System — A group of routers that exchange routing information using the same routing protocol (ex: RIP or OSPF). Packets that are destined for a host outside of the Autonomous System must be sent to the Autonomous System Border router, which in turn sends them to another router in another Autonomous System. (Autonomous System Border routers are described in Section 1.3.1.)

Backbone — Consists of the networks and OSPF routers that do not belong to an area, and routers that belong to multiple areas. OSPF uses the backbone to route packets between areas.

The PowerHub software automatically creates an area ID (0.0.0.0) for the backbone. Unless you add OSPF areas for the interfaces on the PowerHub system, the software assumes that the hub belongs to the backbone area.

Area — A collection of directly-attached OSPF interfaces. The interfaces can be on the same router or on different routers. You assign an OSPF interface to a particular area by assigning a 32-bit area ID to the interface. An interface cannot belong to more than one area.

You can divide your Autonomous System into multiple areas by adding the areas using the **`area add`** command. (See Section 3.3.5 on page 35.)

Stub Area — An area that has only one exit point out of the area.

1.3.1 OSPF Router Types

OSPF recognizes four types of routers. An OSPF router's type is determined by its position in the topology of the OSPF network. Here are the types of OSPF routers:

Backbone router — An OSPF router that has at least one interface attached to the network backbone.

Interior router — An OSPF router whose interfaces all are within the same area.

Area Border router — An OSPF router whose interfaces belong to more than one area.

Autonomous System Border router[1] — An OSPF router that exchanges routing information with routers belonging to other Autonomous Systems.

1. In some documents, the Autonomous System Border router is called the "Autonomous System Boundary Router." Both terms refer to the same type of router.

Figure 1–1 shows an example of an OSPF Autonomous System. In this example, all routers in the Autonomous System belong to the same area (0.0.0.0, the backbone area).

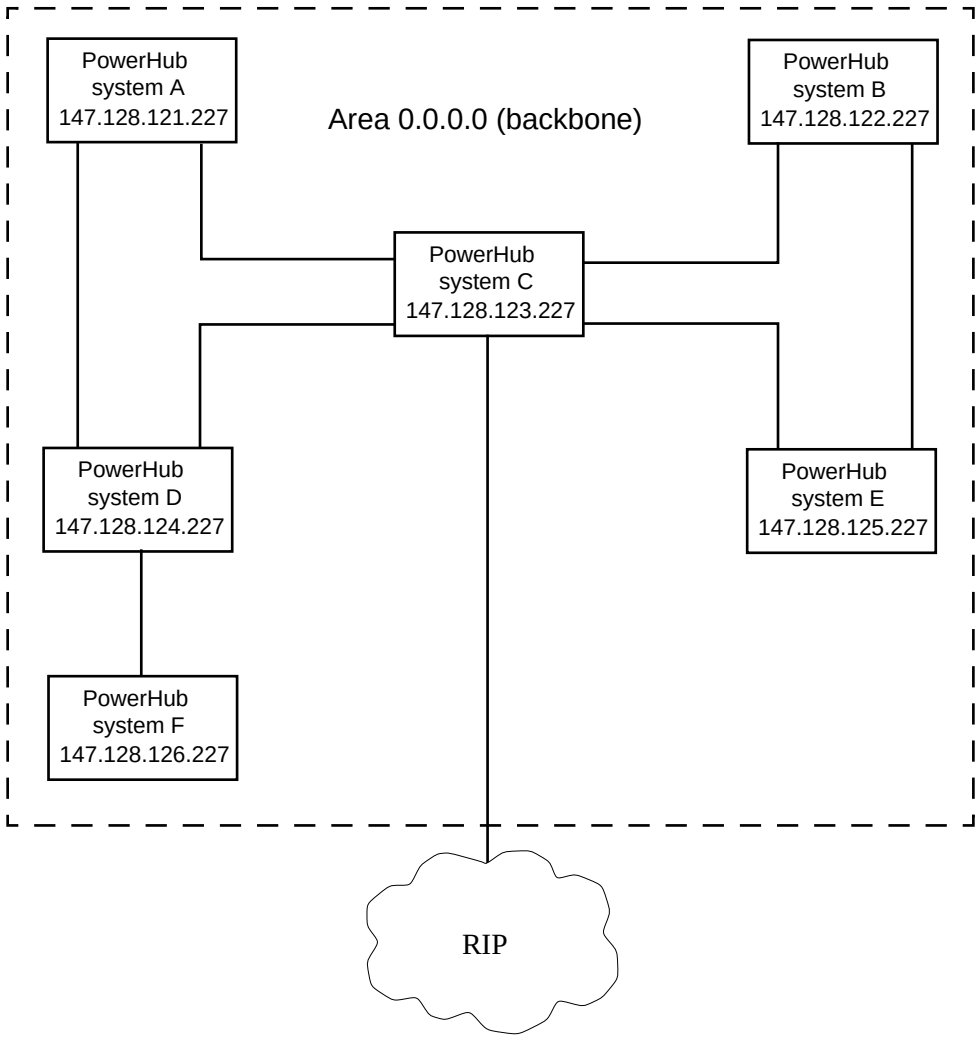

FIGURE 1–1 Example OSPF network containing one area.

The PowerHub systems in Figure 1–1 are functioning as the following types of OSPF routers:

System A	Interior router.
System B	Interior router.
System C	Autonomous System Border router.
System D	Interior router.
System E	Interior router.
System F	Interior router.

Notice that each PowerHub system has an ID that resembles an IP address. These addresses are OSPF router IDs. Each router within an OSPF Autonomous System must have a unique OSPF router ID. To ensure that OSPF IDs remain unique, you can use one of the IP addresses configured on the PowerHub system as that system's OSPF router ID.

Notice that PowerHub system C is attached to another Autonomous System in which RIP is being used. PowerHub system C is functioning as an Autonomous System Border router. The PowerHub software exports OSPF routes into the RIP Autonomous System. The software also imports RIP routes into the OSPF Autonomous System.

Figure 1–2 shows a different configuration using the same PowerHub systems.

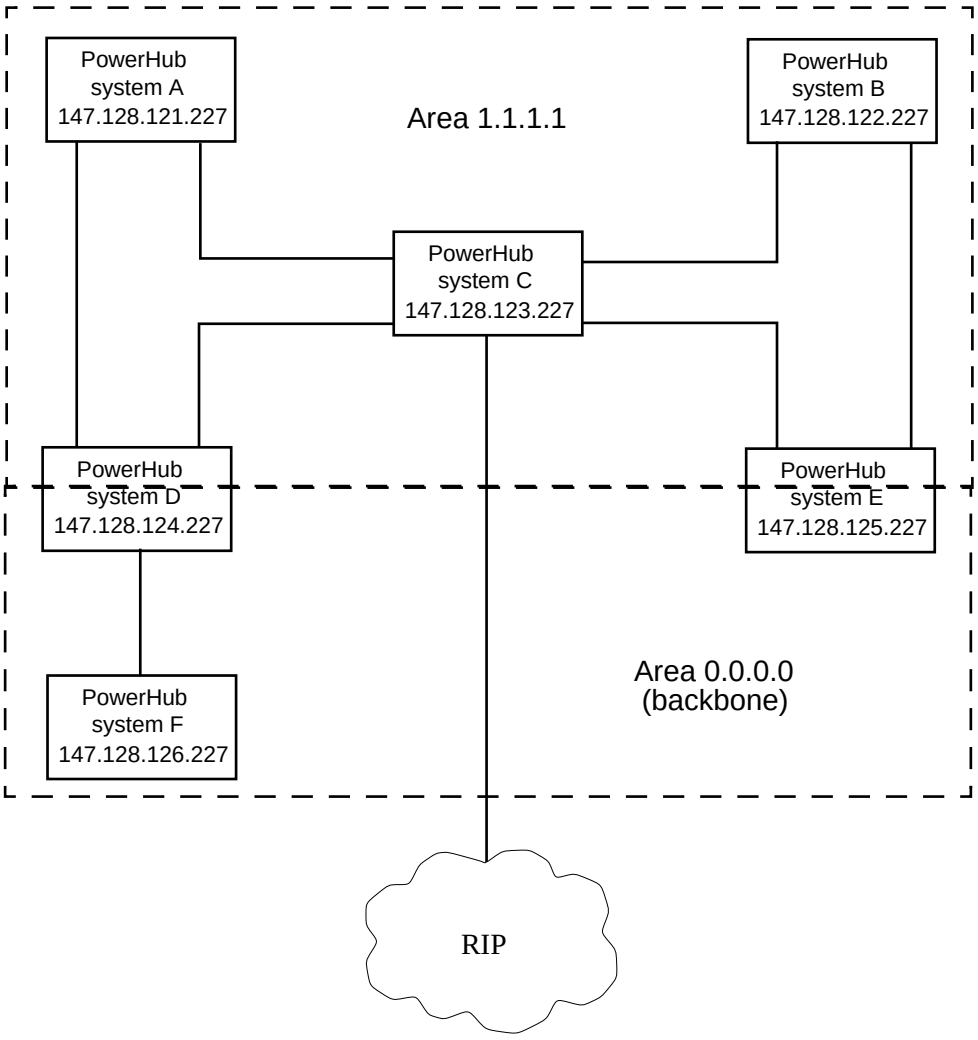

FIGURE 1–2 Example OSPF network containing multiple areas.

In this configuration, the Autonomous System is divided into multiple areas. The PowerHub systems are functioning as the following types of OSPF routers:

System A	Interior router.
System B	Interior router.
System C	Autonomous System Border router.
System D	Area Border router.
System E	Area Border router.
System F	Interior router.

Area Border routers send summary route information from one area to the other. For example, PowerHub systems D and E summarize route information for the networks in area 1.1.1.1 into area 0.0.0.0. In turn, PowerHub systems D and E also summarize route information for area 0.0.0.0 into area 1.1.1.1.

1.3.2 Neighbors

When the OSPF software is initializing the OSPF routers in the Autonomous System, the software uses Hello packets to establish adjacency between routers. Routers that are adjacent exchange their link-state databases to become neighbors. A *link-state database* contains LSAs (link-state advertisements). An *LSA* describes the state of an OSPF router's link to another device within the Autonomous System.

In Figure 1–2 on page 6, PowerHub system A has two OSPF neighbors—PowerHub systems C and D. Following the exchange of Hello packets among these systems, PowerHub system A receives a link-state database from system C and a link-state database from system D. The link-state databases contain link-state information for the networks attached to systems C and D.

OSPF routers use the information in the link-state databases received from their neighbors to route traffic within the Autonomous System. When two neighbors have completely exchanged their link-state databases, they become "full" neighbors. You can display information about the PowerHub system's neighbors using the **`show neighbor`** command. (See Section 3.4.2.4 on page 54.)

1.3.3 Stub Areas

Depending upon your network configuration, you might have an OSPF area that has only one exit. Such an area is called a *stub area*. Figure 1–3 shows an example of an OSPF Autonomous System that contains a stub area.

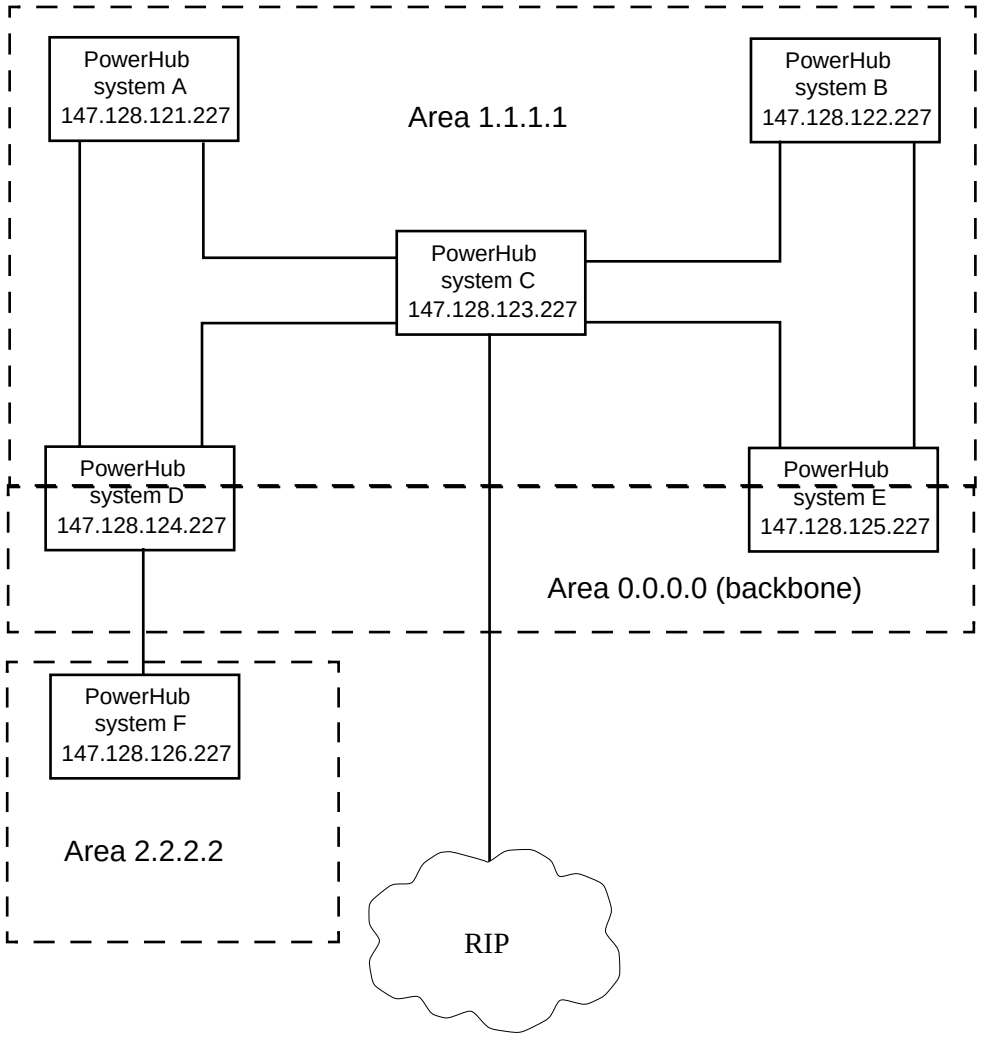

FIGURE 1–3 Example OSPF network containing a stub area.

In this OSPF configuration, PowerHub system F has only one route outside of area 2.2.2.2. All traffic destined for a network inside area 2.2.2.2 must use the single route into the area. All traffic from networks inside area 2.2.2.2 to networks in other areas also must use the single route. The single route into and out of a stub area is called the "default" route for that stub area.

When you add an area to the PowerHub OSPF configuration (using the **`area add`** command), you can explicitly identify the area as a stub area. (See Section 3.3.5 on page 35.) By explicitly identifying the area as a stub area, you permit the PowerHub software to summarize the link-state information sent by other OSPF routers to the stub

area's Area Border router. State information is provided for the default route, rather than for all devices that can be reached using that default route.

1.3.4 Virtual Links

An Autonomous System that contains multiple OSPF areas can contain areas that are not connected to the backbone and are not connected to each other. Figure 1–4 shows an example of an OSPF Autonomous System that contains such an area.

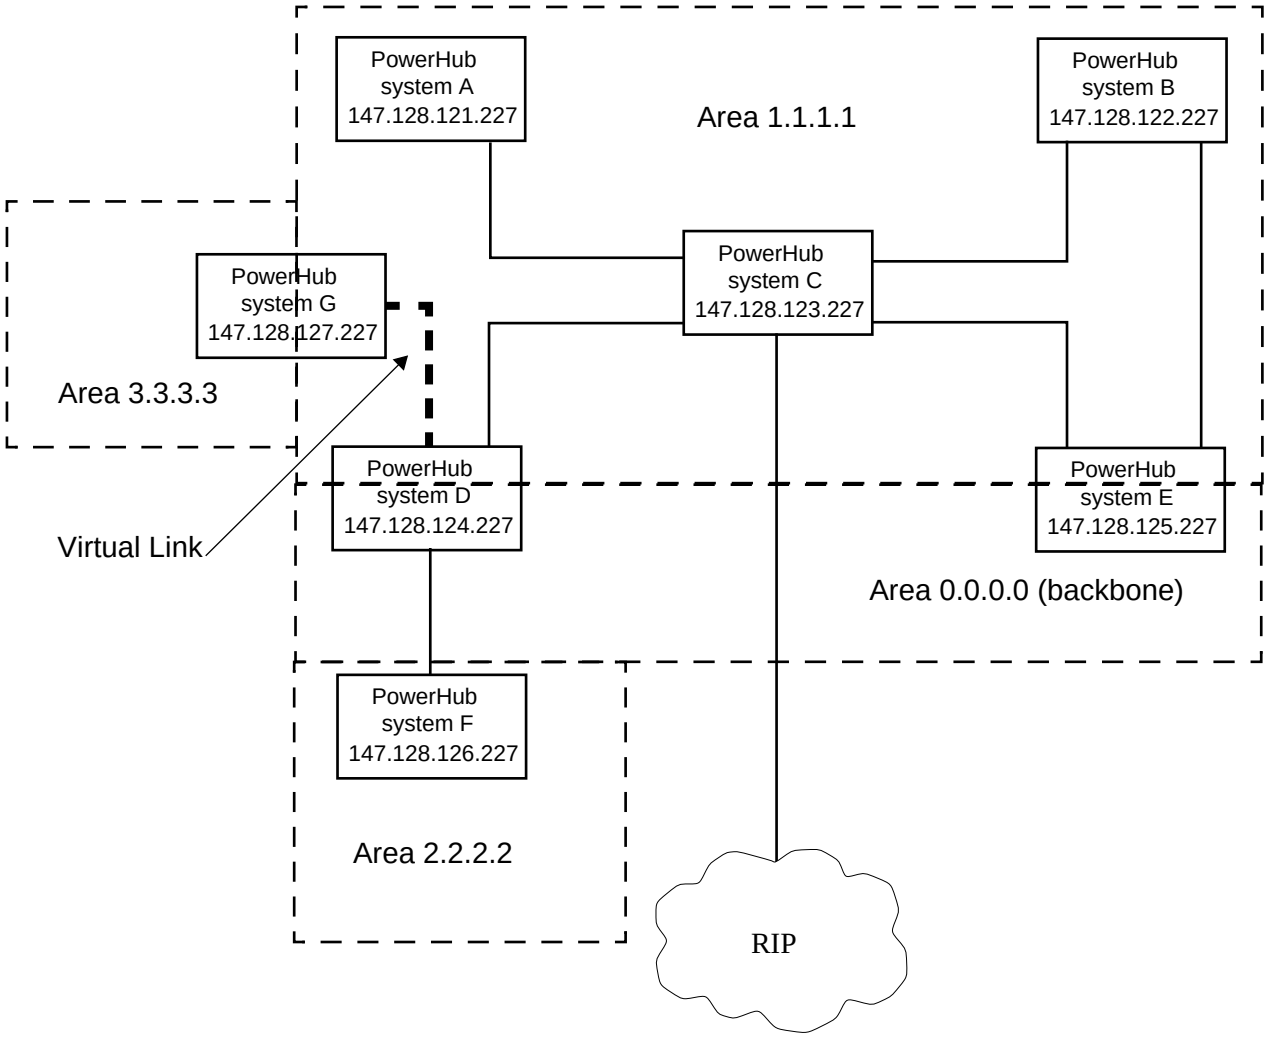

FIGURE 1–4 Example OSPF network containing a virtual link.

In this OSPF configuration, PowerHub system G is disconnected from the backbone area. Normally, routers such as system G would not be able to become neighbors with routers in the backbone area or with areas on the other end of the backbone area. For example, area 3.3.3.3 is disconnected from the backbone area (0.0.0.0) and from area 2.2.2.2. Accordingly, devices inside one of the disconnected areas cannot send traffic to devices in the other disconnected area.

To allow disconnected areas to become neighbors, the PowerHub OSPF software contains the *automatic virtual-link feature*. This feature automatically builds a link between PowerHub OSPF areas that are not attached to the backbone and that do not share a common Area Border router. A virtual link is required to connect system G with a router connected to the backbone, and therefore to the rest of the Autonomous System. The virtual link is shown in Figure 1–4 on page 9 by the heavy dashed line.

The software creates the virtual link by using other areas between the disconnected areas as transit areas. A *transit area* provides a path for the virtual link. In Figure 1–4, the transit area for the virtual link between system G and the backbone area is area 1.1.1.1.

Note that if one or both of the disconnected OSPF routers is not a PowerHub system, the PowerHub OSPF software cannot automatically build a virtual link between the routers. However, you still can build a virtual link between the routers using the **`virtual-link add`** command. (See Section 3.3.8 on page 41.)

1.3.5 Exchanging Routes Between RIP and OSPF

You can use both OSPF and RIP on the same PowerHub system, as long as you don't use both routing protocols on the same PowerHub segment. A PowerHub system running both RIP and OSPF can exchange route information between the two protocols.

A PowerHub system that runs both OSPF and RIP can serve as an Autonomous System Border router. As described in Section 1.3.1 on page 4, an Autonomous System Border router connects an OSPF Autonomous System to another Autonomous System where another routing protocol, such as RIP, is being used. (In Figure 1–4, PowerHub system C is an Autonomous System Border router.)

If you want the hub to exchange RIP and OSPF information, you must enable the hub as an Autonomous System Border router using the **`set asbd enl`** command. (See Section 3.3.2 on page 34.)

2 Getting Started

This chapter contains procedures for configuring your PowerHub system to perform OSPF routing. The procedures assume that you want to use the PowerHub defaults for most of the OSPF parameters. Use the procedures in this chapter to perform the basic configuration tasks, then use specific **ospf** commands described in Chapter 3 to fine-tune your configuration if needed.

This chapter contains the following information:

- Displaying the installed system software version. (See Section 2.1 on page 11.)
- Displaying the amount of DRAM in the Packet Engine. (See Section 2.2 on page 12.)
- Accessing the **ospf** subsystem. (See Section 2.3 on page 13.)
- Getting on-line help. (See Section 2.4 on page 13.)
- Configuring the PowerHub system for OSPF routing. (See Section 2.5 on page 14.)
- Saving the PowerHub configuration. (See Section 2.6 on page 21.)
- Verifying and troubleshooting the OSPF configuration. (See Section 2.7 on page 24.)

2.1 DISPLAYING THE SYSTEM SOFTWARE VERSION

Section 1.1 on page 2 lists the hardware and software requirements for using the PowerHub OSPF software. To display the versions of the system software and Packet Engine boot PROM installed on your PowerHub system, issue the following command:

`main version` *`<slot-number>`*

where:

`<slot-number>` Is the number of the NIM slot that contains the Packet Engine. Specify one of the following:

- For a PowerHub 4000, 4100, or 6000, specify **`1`**.
- For a PowerHub 7000, specify **`5`** for a 5-slot chassis. Specify **`10`** for a 10-slot chassis.

Here is an example of the information shown by this command. In this example, information is shown for a 5-slot PowerHub 7000.

```
1:PowerHub:main# version 5
############## Slot 5 ##############

       Card Type: Packet Engine
        Serial #: 538027795
           Model: 7101-00
        Revision: G
           Issue: 1

PowerHub Version: 7-2.6.3.0 (s1.91) 1995.12.07 14:02
    PROM Version: 7pep-2.5.4 (s1.76) 1995.11.30 15:41
```

The numbers in parentheses following the software version names are used by Technical Support and might differ from the number shown in this example.

Notice that the date and time when the final versions of the software were officially released by the factory are listed to the right of the software versions. (The release dates and times shown in this example might differ from those actually displayed on your PowerHub system.)

2.2 DISPLAYING THE DRAM IN THE PACKET ENGINE

To ensure that your PowerHub system contains enough DRAM to run OSPF, issue the **mgmt showcfg** command. This command displays the amount of DRAM in your Packet Engine. Here is an example of the display produced by the **mgmt showcfg** command. The line shown in bold type lists how much DRAM is in the Packet Engine.

```
1:PowerHub:main# mgmt showcfg
Accelerator board is present. Accelerator IOP is being used.
Installed DRAM Size: 24 MB
tty1 not set - using 9600 baud
tty2 not set - using 1200 baud
PE: slot 5

04/19: FDDI    FDDI
03/13: FIBER   FIBER   FIBER   FIBER   FIBER   FIBER
02/14: BNC     BNC     BNC     BNCT    BNCT    BNCT
01/01: UTP     UTP     UTP     UTP     UTP     UTP     UTP     UTP
       UTP     UTP     UTP     UTP     100TX
```

In this example, the Packet Engine contains a 24 MB SIMM.

2.3 ACCESSING THE OSPF SUBSYSTEM

Most of the commands described in this chapter are located in the **ospf** subsystem. To access the **ospf** subsystem, issue the following command at the runtime prompt:

`ospf`

The procedures in this chapter assume that you will be working from within the **ospf** subsystem. Accordingly, commands in the **ospf** subsystem are not shown with the subsystem name. However, for commands that are not located in the **ospf** subsystem, the subsystem name is included as part of the command. For example, the **`add-interface`** command in the **ip** subsystem is shown as **`ip add-interface`**.

2.4 GETTING ON-LINE HELP

All PowerHub models have on-line help to help you navigate through the user interface. You can get the following types of on-line help:

- General help.
- Subsystem help.
- Command-specific help.

To access general help, issue the following command:

`findcmd|fcmd`

This command lists all the commands in all the subsystems.

To display a list of commands within a subsystem, issue the following command:

`help|h`

This command displays information for only the subsystem which you are currently using.

To use command-specific help, issue the following command:

`help|h` *<command-name>*

Command-specific help provides information for only the command you specify. Separate "**`help`**" from the command name by using a space.

2.5 CONFIGURING THE HUB AS AN OSPF ROUTER

As described in Section 1.3.1 on page 4, you can configure the PowerHub system as the following types of OSPF router:

- Interior
- Backbone[1]
- Area Border
- Autonomous System Border

An OSPF router can function as more than one of the router types listed above. For example, a PowerHub system that has interfaces attached to the backbone and to other OSPF areas will function both as a Backbone router and as an Area Border router.

Generally, you do not need to worry about the differences among these router types. The PowerHub OSPF software determines how the hub is being used based upon your network configuration.

To configure the PowerHub system for OSPF routing, you need to perform the following tasks. These tasks apply to all OSPF router types.

- Allocate memory for OSPF.
- Add IP interfaces (if interfaces are not already configured).
- Enable IP forwarding (if it is not already enabled).
- Assign the OSPF router ID.

Depending upon the type of OSPF router you plan to use the PowerHub system as, you might need to perform some additional configuration tasks.

- If you plan to use the hub as an Area Border router, you need to add OSPF areas, then add OSPF interfaces to the areas.
- If your network contains areas that are not connected to the backbone and are not connected to each other, and the Area Border router for one of these areas is not a PowerHub system, you might need to create virtual links.
- If you plan to use the hub as an Autonomous System Border router, you need to enable the hub as this type of router.

Finally, after you complete the OSPF configuration steps listed above, you need to enable OSPF routing.

The following sections describe how to perform these tasks.

---

1. You do not need to perform any special configuration steps to use the PowerHub system as a Backbone router. If your configuration of the hub as an Area Border router or Autonomous System Border router places the hub on the backbone, the hub also functions as a Backbone router. An Interior router cannot be a Backbone router.

   Unless you configure OSPF areas using the **`area add`** command, the PowerHub software assumes that you are configuring the PowerHub system as a Backbone router. In addition, the software automatically configures the area ID 0.0.0.0 for the backbone.

2.5.1 PowerHub OSPF Defaults

The procedures in this chapter assume that you are using the OSPF configuration defaults listed in Table 2–1.

TABLE 2–1 PowerHub OSPF defaults.

OSPF Parameter	PowerHub Default
OSPF routing	Disabled
Interface state	Enabled (when you enable OSPF routing)
Area	0.0.0.0 (backbone)
Network ranges	None. Link-state information for all networks attached to the hub is reported in hub's link-state database.
Autonomous System Border router (enables hub to exchange RIP and OSPF routes)	Disabled
Automatic virtual-link feature	Enabled

If your network configuration requires that you use different values for any of these parameters, first complete the setup procedures in this chapter. Then use commands in the **ospf** subsystem to adjust the OSPF parameters on the PowerHub system. For information about all the commands in the **ospf** subsystem, see Chapter 3.

2.5.2 Allocating Memory

You must allocate a portion of the PowerHub system's main memory for the **ospf** subsystem. If you do not allocate memory to the **ospf** subsystem, you cannot access it.

To allocate memory for the **ospf** subsystem, issue the following command:

`main getmem ospf [`*`<Kbytes>`***`]`**

where:

`<Kbytes>` Specifies how many KB of the Packet Engine's main memory you want to allocate to OSPF. You can specify an amount from **`10`** through **`500`** KB, in 1-KB increments. The default is **`500`** KB.

FORE Systems recommends that you try the default allocation first. Let the PowerHub system act as an OSPF router in your OSPF network for a few hours, then use the **`stats`** command to check the memory usage. If the `Free` field shows more than 25% of the memory is unused, you might want to change the memory allocation to a smaller amount. (See Section 3.5.1 on page 60.)

In general, you want to allocate enough memory so that about 25% is free.

If the **`stats`** command indicates memory allocation failures (in the `Total Authorization Failures` field), you need to allocate more memory to OSPF.

NOTE: You can allocate memory for a specific protocol (such as OSPF) only one time during a power cycle. After the **main getmem ospf** command is issued once within a power cycle (either from within the PowerHub configuration file or by you), the command does not affect the amount of DRAM allocated to OSPF.

To change the amount of DRAM allocated to OSPF, issue the **main getmem ospf** command, specifying the amount of DRAM you want to allocate, save the configuration file, then reboot the hub. When you reboot, the hub reads the new DRAM amount from the configuration file. See Section 2.6 on page 21 for information about saving the configuration file.

2.5.3 Adding an IP Interface

When you add an IP interface to the PowerHub system, the PowerHub software automatically adds a corresponding OSPF interface. The OSPF interface is enabled when you enable OSPF routing.

If you are adding OSPF to a PowerHub system that is already configured for your IP network, skip the rest of this section and go to Section 2.5.5 on page 18. The OSPF software automatically creates OSPF interfaces for the IP interfaces when you enable OSPF routing.

If you need to configure IP interfaces on the PowerHub system, use the following command.

ip add-interface|ai *<seg-list> <IP-addr>*
[*<subnet-mask>* **[br0|br1]**
[cost *<cost>***] [allsubnets|as]]**

where:

<seg-list> — Specifies the segment(s) to which you are assigning the IP address. You can specify a single segment, a comma-separated list of segments, a hyphen-separated range of segments, or **all** for all segments. If you specify multiple segments, you create a VLAN. (See Appendix D in the *PowerHub Software Manual, V 2.6* for information about VLANs.)

<IP-addr> — Specifies the IP address you want to assign to the specified segment(s). The IP address must be in dotted-decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9).

<subnet-mask> — Specifies the subnet mask. If a particular network uses IP subnet addressing, then the subnet mask should be specified here using dotted-decimal notation. Otherwise, the system uses a default subnet mask equal to the "natural" subnet mask for the particular class of address.

br0|br1 — Specifies the style of broadcast address on a segment-by-segment basis:

- If you specify **br0**, the hub sends an "all-0s" broadcast. This means all bits in the host segment of the address are 0s.
- If you specify **br1**, the hub sends a standard "all-1s" broadcast. This means all bits in the host segment of the address are 1s. The default is **br1**.

cost *<cost>* — Specifies an additional cost of using the subnet interface. This is the number of extra hops to the destination. The range is **1** through **14**. (The router decrements an IP packet's time-to-live field at each hop.) The default is **0** (zero). When the hub reports this subnet using RIP, it adds the additional cost to the reported metric.

NOTE: The cost you specify using the **cost** argument is used only by RIP, not by OSPF.

allsubnets|as — Lets you specify all zeros or all ones in the subnet mask. You must use this argument if you specify either of the following subnet masks: **0.0.0.0** or **255.255.255.255**.

For information about the **ip add-interface** command and other commands in the **ip** subsystem, see Chapter 5 in the *PowerHub Software Manual, V 2.6* (Rev B).

2.5.4 Enabling IP Forwarding

If you are adding OSPF to a PowerHub system that is already configured for your IP network, IP forwarding probably is already enabled. To verify whether IP forwarding is enabled, issue the following command:

```
ip showcfg
```

Here is an example of the display produced by this command.

```
1:PowerHub:ip# showcfg
IP Configuration:
----------------
IP Forwarding:                      enabled (gateway)
Default TTL:                        32
Arp cache aging time:               5 minutes
Routing Network Broadcasts:         enabled
VLAN Bridging Network Broadcasts:   enabled
Send ICMP redirects:                enabled
Forward Pkts with SrcRt Option:     enabled
Routed Packet Snooping:             disabled
```

The `IP Forwarding` field shows whether IP forwarding is enabled. (This line is shown in bold type in the example.)

If IP forwarding is enabled, go to Section 2.5.5.

If IP forwarding is disabled, enable it using the following command:

```
ip set ipForwarding enl
```

2.5.5 Assigning the OSPF Router ID

Each OSPF router within the Autonomous System must have a unique OSPF router ID. The OSPF router ID is a 32-bit address in IP format. The software does not assign an address automatically.

You can use any 32-bit address for the OSPF router ID. However, FORE Systems recommends that you use one of the IP addresses configured on the PowerHub system. By using one of the IP address on the PowerHub system, you can ensure that OSPF IDs remain unique. If you choose an IP address configured on the hub, your choice does not affect IP or OSPF. That is, the software does not establish a special relationship between the IP address you choose and the OSPF software.

By requiring that you use an IP address configured on the hub, the PowerHub OSPF software ensures that your OSPF router ID remains unique regardless of changes in the network.

To assign the OSPF router ID, issue the following command:

set|se router-id|ri *<router-id>*

where:

<router-id> Specifies the OSPF router ID. Specify the router ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 through 9).

NOTES: You can define the OSPF router ID only when OSPF routing is disabled. To verify that OSPF routing is disabled, issue the **showcfg** command (see Section 2.7.1 on page 24).

Here is an example of this command.

```
2:PowerHub:ospf# set router-id 1.1.1.1
```

2.5.6 Additional Configuration Steps

Unless you plan to use the PowerHub system as an Interior router, you need to perform some additional configuration steps:

- If you plan to use the PowerHub system as an Area Border router, you need to add OSPF areas, then add OSPF interfaces to those areas. (See Section 2.5.6.1 and Section 2.5.6.2.) You might also need to add areas and interfaces if you plan to use the hub as an Autonomous System Border router.
- If the PowerHub system is in an area that is disconnected from the backbone and from another area, and the Area Border router for the other area is not a PowerHub system, you might need to manually add a virtual link. (See Section 2.5.6.3.)
- If you plan to use the PowerHub system as an Autonomous System Border router, you need to enable the hub to be this type of router. (See Section 2.5.6.4.)

2.5.6.1 Adding an OSPF Area

The PowerHub OSPF software automatically creates an area (0.0.0.0) for the network backbone. Unless you explicitly add the OSPF interfaces on the PowerHub system to an area, the software configures the interfaces to belong to the backbone area (0.0.0.0). After adding an area, you can add OSPF interfaces to the area using the **interface** command. (See Section 2.5.6.2.)

To add an area to the PowerHub system, issue the following command:

area|ar add|a *<area-id>* **[***<auth-type>***] [stub-area|sa** *<cost>***]**

See Section 3.3.5 on page 35 for information about the command syntax.

Here is an example of this command. In this example, an area with ID 3.3.3.3 is added. The *<auth-type>* argument is used to specify that a simple password is required for this area.

```
3:PowerHub:ospf# area add 3.3.3.3 sp
```

2.5.6.2 Adding an OSPF Interface to an Area

When you enable OSPF routing (see Section 2.5.7 on page 21), the PowerHub OSPF software automatically enables all OSPF interfaces and assigns them to the backbone area (0.0.0.0). You can assign an interface to a different area, disable the interface, or change the values of other interface parameters using the following command:

interface|if *<ip-addr>* **[ar** *<area-id>***] [auth** *<key-str>***]**
[cost|c *<cost>***] [priority|p** *<priority>***]**
[xdelay|x *<trans-delay>***] [rint|r** *<rxmt-int>***]**
[hint|h *<hello-int>***] [rdint|d** *<rtr-dead-int>***]**
[state|s *<state>***]**

See Section 3.3.6 on page 37 for information about the command syntax.

Here is an example of this command. In this example, interface 192.9.222.1 is added to area 3.3.3.3. The password “MKT” is used for authentication. (In the example in Section 2.5.6.1 on page 19, area 3.3.3.3 was configured to use a simple password for authentication.)

```
4:PowerHub:ospf# interface 192.9.222.1 ar 3.3.3.3 auth MKT
OSPF: Interface 192.9.222.1 defined
```

2.5.6.3 Adding a Virtual Link

If your Autonomous System contains more than a single OSPF area, it is possible for some areas to be unable to reach other areas. This situation can occur when two or more areas do not have interfaces in the backbone area and do not share an Area Border router.

The PowerHub OSPF software contains the automatic virtual-link feature. This feature automatically creates links between areas that cannot reach each other. If the Area Border routers in both areas are PowerHub systems, the feature can build the virtual links.

However, the automatic virtual-link feature establishes virtual links only between PowerHub systems. If the Area Border router in an OSPF area is not a PowerHub system, you need to manually create the virtual link between that area and an area that uses a PowerHub system as the Area Border router.

To build a virtual link manually, issue the following command:

```
virtual-link|vl add|a <aid> <router-id> [auth <key-str>]
                      [xdelay|x <trans-delay>]
                      [rint|r <rxmt-int>]
                      [hint|h <hello-int>]
                      [rdint|d <rtr-dead-int>]
```

See Section 3.3.8 on page 41 for information about the command syntax.

NOTE: You must perform this step before you enable OSPF routing.

2.5.6.4 Autonomous System Border Router

If you are using the PowerHub system to join two Autonomous Systems, you must enable the hub to be an Autonomous System Border router. Enabling the hub as an Autonomous System Border router allows the hub to exchange route information between OSPF and RIP.

To enable the hub as an Autonomous System Border router, issue the following command:

```
set asbd enl
```

NOTE: You must perform this step before you enable OSPF routing.

2.5.7 Enabling OSPF Routing

After you complete the applicable configuration tasks in the previous procedures, you are ready to enable OSPF routing. To enable OSPF routing, issue the following command:

```
set OSPF-router enl
```

Here is an example of the system response when you enable OSPF routing.

```
5:PowerHub:ospf#  set ospfr enl
OSPF: Routing Enabled
```

2.6 SAVING THE POWERHUB CONFIGURATION

After you finish configuring the PowerHub system for OSPF routing, you should save the system configuration. You save the system configuration by saving the PowerHub system's configuration file. When you boot the hub, the hub reads the configuration file and configures the hub according to the contents of the file.

NOTE: If you plan to use more than one boot source, we recommend that you issue the commands for saving the configuration onto each boot source. These procedures are described in the sections below. You can avoid potential problems by always ensuring that the configuration files on all boot sources match.

2.6.1 Floppy Drive

To save the system configuration onto the diskette in the floppy drive:

(1) Boot the software, if you have not already done so. Following the boot messages, a command prompt similar to the following is displayed:

```
1:PowerHub:
```

NOTE: The software must be successfully booted before you can save the configuration. You cannot save the configuration from the boot-PROM prompt.

(2) Type the following command at the prompt, then press Enter:

```
mgmt savecfg fd:cfg
```

This command saves the configuration into a file named `cfg` on the floppy diskette. You can specify any DOS-like file name (`xxxxxxxx.xxx`). If you specify a name that does not match the name in the `bootdef` file (`cfg`), you must load the configuration manually each time you boot the software by using the **`mgmt readcfg`** command.

2.6.2 Flash Memory Module

To save the system configuration onto the Flash Memory Module:

(1) Boot the software, if you have not already done so. Following the boot messages, a command prompt similar to the following is displayed:

```
1:PowerHub:
```

NOTE: The software must be successfully booted before you can save the configuration. You cannot save the configuration from the boot-PROM prompt.

(2) Type the following command at the prompt, then press Enter:

`mgmt savecfg fm:cfg`

This command saves the configuration into a file named `cfg` on the Flash Memory Module. You can specify any DOS-like file name (`xxxxxxxx.xxx`). If you specify a name that does not match the name in the `bootdef` file (`cfg`), you must load the configuration manually each time you boot the software by using the **`mgmt readcfg`** command.

2.6.3 TFTP Server

To save the system configuration file onto a TFTP server, you can do either of the following:

- Save the configuration file to the floppy drive using the procedure in Section 2.6.1 on page 21, then place the diskette into a disk drive attached to your TFTP server and copy the file from the diskette to the server.
- Use the PowerHub TFTP software to directly save the configuration into a file on your TFTP server. This section describes how to do this.

NOTE: The segment that connects the PowerHub system to the TFTP server must have an IP interface. The TFTP software uses the IP interface.

(1) Boot the software, if you have not already done so. Following the boot messages, a command prompt similar to the following should be displayed:

```
1:PowerHub:
```

NOTE: The software must be successfully booted before you can save the configuration. You cannot save the configuration from the boot-PROM prompt.

(2) Attach a segment cable from the PowerHub system to the TFTP server if a cable does not already connect them.

(3) Type the following command at the prompt, then press Enter:

`tftp svcfg [-h` *`<host>`***`]`** *`<remfile>`*

where:

`-h` *`<host>`* Specifies the IP address of the TFTP server. Unless you have already specified a default TFTP server using the **`tftp set`** command, you need to include this argument. For information on the **`tftp set`** command, see the *PowerHub Software Manual, V 2.6*.

`<remfile>` Specifies the configuration file name. Specify a name that is meaningful to the TFTP program on the server. For example, if the server contains a subdirectory called `fore` and this directory is specified as the TFTP home directory, do not specify `fore` as part of the file name.

NOTES: Some TFTP servers require that the remote file name exist on the server before you can write to that file name. If your server requires that the file name already exist, create a short file (named the same as your configuration file) on the server, then specify that file name for *`<remfile>`*.

Also, on some TFTP servers, including servers running Sun/OS 4.x, files that you overwrite on the server are not properly truncated. When you overwrite an existing file on the TFTP server, if the older version of the file is longer than the new file, the older version is not truncated properly by the server. As a result, the new version of the file contains part of the older version of the file.

If the configuration file name you specified in the boot definition file on the server is longer than eight characters, you can copy the file to the server by using a DOS-like name, then rename the file on the server to match the file name you specified in the boot definition file.

2.7 VERIFYING THE OSPF CONFIGURATION

After you complete the configuration tasks in this chapter, you can verify the OSPF configuration using the following commands:

showcfg — Displays basic information such as whether memory has been allocated for OSPF and whether OSPF forwarding is enabled.

show *<ospf-parm>* — Displays detailed information about specific OSPF configuration items such as areas and interfaces.

2.7.1 Verifying Basic Configuration Settings

To display the hub's OSPF configuration, issue the following command:

showcfg|scf

This command indicates whether memory has been allocated for OSPF and indicates the state (enabled or disabled) of OSPF features, as shown in the following example.

```
2:PowerHub:ospf# showcfg
OSPF Router                          : Memory Available
OSPF Routing                         : Enabled
OSPF Router ID                       : 1.1.1.1
OSPF Version Number                  : 2
OSPF Autonomous System Border Router : Enabled
Automatic Virtual Link Feature       : Enabled
```

The fields in this display show the following information:

`OSPF Router` — Indicates whether you have allocated a portion of the PowerHub memory for OSPF. In this example, memory has been allocated. If memory has not been allocated, you can allocate the memory using the following command: **main getmem ospf**. (See Section 2.5.2 on page 15.)

`OSPF Routing` — Indicates whether OSPF routing is enabled. If OSPF routing is disabled, enable it using the **set OSPF-routing enl** command. (See Section 2.5.7 on page 21.)

`OSPF Router ID` — Displays the OSPF router ID you assigned to this PowerHub system. If no router ID is displayed, assign the router ID using the **set router-id** command. (See Section 2.5.5 on page 18.)

`OSPF Version number` — Indicates that OSPF version 2 is implemented on the PowerHub system. This value does not change.

`OSPF Autonomous System Border Router` — Indicates whether you have configured the PowerHub system as an Autonomous System Border router. This feature is disabled by default, but you should enable the feature if you want the hub to exchange routes between RIP and OSPF. To enable the feature, issue the **set asbd enl** command. (See Section 3.3.3 on page 35.)

`Automatic Virtual Link Feature`

Indicates whether the automatic virtual-link feature is enabled. This feature is enabled by default. FORE Systems recommends that you leave this feature enabled at all times. For information about enabling the feature, see Section 3.3.3 on page 35.

2.7.2 Verifying Communication with Neighbors

To perform OSPF routing, the PowerHub system must exchange its link-state database with the OSPF routers that are attached to the hub. The OSPF routers attached to the hub are the hub's OSPF neighbors.

To display the status of the hub's relationships with its neighbors, issue the following command:

`show neighbor`

Here is an example of this command. In this example, the PowerHub system has two neighbors.

```
3:hub2:ospf# show neighbor
IP Address      Router ID       Pri State    Events RTrQ
--------------- --------------- --- -------- ------ -----
129.213.72.2    5.5.5.5         1   full     6      0
150.1.100.3     3.3.3.3         1   full     6      0
```

The `State` field lists the state of the hub's relationship with this neighbor. In this example, the hub and its neighbor have completely exchanged their link-state databases. By exchanging their link-state databases, the hub and its neighbor each have obtained the route information known by the other.

If the `State` field contains a value other than "`full`" (for example, `attempt`), one of the following conditions might be true:

- The hub and its neighbor are still exchanging their link-state databases.
- The hub and its neighbor are using different authentication types for the same area.
- The interface connecting the hub and its neighbor are using different authentication strings, even if the hub and its neighbor are using the same authentication type for the area.
- The Hello protocol failed to establish adjacency between the hub and its neighbor.

Wait about 30 seconds for the hub and its neighbor to finish exchanging their link-state databases, then issue the **`show neighbor`** command again. If the state still is not "`full`," use the methods in the following sections to diagnose the problem.

See Section 3.4.2.4 on page 54 for a description of all the possible neighbor states.

2.7.2.1 Verifying Interface States

All OSPF interfaces are enabled by default. However, an OSPF interface can become disabled in one of the following ways:

- You disable it using the **`interface`** *<ip-addr>* **`state disable`** command.
- You disable it using the **`bridge port`** command.
- The PowerHub software's automatic segment-state detection feature disables it.

To verify the state of the OSPF interfaces, issue the following command:

`show interface`

Here is an example of the information displayed by this command. The interface state is listed in the `Admin` field.

```
4:PowerHub:ospf# show interface
IP Address      Area Id         DR              BDR             Admin
--------------- --------------- --------------- --------------- ---------
192.9.222.1     3.3.3.3         0.0.0.0         0.0.0.0         Enabled
150.1.200.1     0.0.0.0         0.0.0.0         0.0.0.0         Enabled
129.213.72.1    0.0.0.0         129.213.72.2    129.213.72.1    Enabled
150.1.100.1     1.1.1.1         150.1.100.3     150.1.100.1     Enabled
```

If an interface is disabled, check the link LED for the segment. If the LED is dark, check the segment cable to ensure that the cable is undamaged and is connected at both ends.

If the cable is attached and the link LED is lit, the OSPF interface might be disabled. Issue the following command to enable the OSPF interface:

`interface` *<ip-addr>* **`state enable`**

2.7.2.2 Verifying the Authentication Type and String

The authentication type is associated with the area. The authentication string is associated with the interface.

To check the authentication type configured on the PowerHub system for an area, issue the following command:

`show area` *<aid>*

where:

<aid> Is the ID of the area shared by the hub and its neighbor.

Here is an example of the information displayed by the **show area** command. The authentication type is listed in the Auth Type field.

```
5:PowerHub:ospf# show area
                Auth Import AS Number of  # Area # AS   Number of  Stub Area
Area Id         Type Ext LSAs  Spf Runs   Bdr    Bdr    Area LSAs  Cost
--------------- ---- --------- ---------- ------ ------ ---------- --------
0.0.0.0          no  Enabled   12         4      4      13         -----
1.1.1.1          no  Enabled   12         2      2      15         -----
1.2.3.4          md5 Enabled   12         0      0      0          -----
2.3.4.5          sp  Enabled   12         0      0      0          -----
3.3.3.3          sp  Enabled   12         1      1      16         -----
33.0.33.0        no  Enabled   12         0      0      0          -----
33.33.33.33      no  Enabled   12         0      0      0          -----
```

The Auth Type field can contain one of the following values:

no	No authentication is required for this area.
sp	A simple password is required for this area. You specify the password when you create the area using the **area add** command.
md5	MD5 authentication is required in this area. See RFC 1321 for information about MD5.

If the neighbor also is a PowerHub system, issue this command on the neighbor too. If the neighbor is not a PowerHub system, see the documentation for the neighbor for information about displaying the authentication type used on the neighbor.

To verify the authentication string defined for an interface, issue the following command:

show interface *<ip-addr>*

where:

<ip-addr>	Is the OSPF interface address.

The authentication string for the interface is displayed in the Authorization Key String field (shown in bold type in this example).

```
6:hub2:ospf# show interface 192.9.222.1
168:ospf-hub1:ospf# show interface 192.9.222.1
IP Address                  : 192.9.222.1
Area ID                     : 3.3.3.3
Interface Type              : Broadcast
Administrative Status       : Enabled
Router Priority             : 1
Interface Metric            : 10
Transit Delay               : 1
Retransmission Interval     : 5
Hello Interval              : 10
Router Dead Interval        : 40
Interface State             : Down
Designated Router           : 0.0.0.0
Backup Designated Router    : 0.0.0.0
Interface Events            : 0
Authorization Key String    : MKT
Authorization Failures      : 0
```

2.7.2.3 Verifying that the Hub and Neighbor Are Exchanging Packets

If the commands described in the previous sections do not indicate an error, but the PowerHub system still is not performing OSPF routing, display OSPF statistics to look for packet errors or memory allocation errors.

To display link-state database statistics, memory usage statistics, and other OSPF statistics for the PowerHub system, issue the following command:

`stats`

Here is an example of the information displayed by this command.

```
7:hub2:ospf# stats
External Link-State Advertisements        : 0
Checksum of the External LSA Database     : 0
New Link-State Advertisements originated  : 105
Link-State Advertisements received        : 121
OSPF Area Border Router                   : True
Total Authorization Failures              : 0
ospf: Total 100 KB. Used 33 KB. Free  66 %
```

If the value is not 0 (zero) for the `Total Authorization Failures` field (shown in bold type in this example), you need to allocate a larger amount of DRAM to OSPF. (See Section 2.5.2 on page 15.)

2.7.3 Verifying the OSPF Routes

The IP route table lists the routes known to the PowerHub system. You can display the IP route table to verify the presence of routes to the other networks within the autonomous system.

To display the IP route table, issue the following command:

ip route-table|rt
[-c|-d|-r|-s] [-f] [-t] [<*seg-list*>**] [**<*IP-addr*>**]**

Here is an example of the information displayed by this command. See Section 5.11 in the *PowerHub Software Manual, V 2.6* (Rev B) for information about the command syntax.

```
8:PowerHub:ip# route-table

  Destination     Subnet Mask       Gateway     Met State RtSrc  Age    Port
--------------- --------------- --------------- --- ----- ------ ---- ---------
44.0.0.0        255.0.0.0       150.1.100.3      30 Activ OIntra ---- 14
80.100.0.0      255.255.0.0     150.1.100.3      20 Activ OIntra ---- 14
80.200.0.0      255.255.0.0     150.1.100.3      30 Activ OIntra ---- 14
87.0.0.0        255.0.0.0       150.1.100.3      30 Activ OInter ---- 14
111.222.111.0   255.255.255.0   ---------------   0 Down  Direct ---- none
150.1.100.0     255.255.255.0   ---------------   0 Activ Direct ---- 14
150.1.200.0     255.255.255.0   ---------------   0 Down  Direct ---- none
152.16.0.0      255.255.0.0     150.1.100.3      30 Activ OInter ---- 14
192.9.222.0     255.255.255.0   ---------------   0 Down  Direct ---- none
213.100.45.0    255.255.255.0   150.1.100.3      20 Activ OIntra ---- 14
170.170.1.0     255.255.255.0   150.1.100.3      20 Activ OInter ---- 14
180.180.0.0     255.255.0.0     150.1.100.3      30 Activ OIntra ---- 14
129.213.0.0     255.255.0.0     ---------------   0 Activ Direct ---- 19
200.200.200.0   255.255.255.0   150.1.100.3      30 Activ OIntra ---- 14

Total routes: 14 (Direct: 5, Static: 0, RIP: 0, OSPF: 9)
Up routes: 11, Down routes 3
```

Examine the routes listed in the IP route table to ensure that all the routes you expect the hub to know about are present in the table.

For each IP route, the route table shows the following information:

`Destination` The IP address of the destination host or net.

`Subnet Mask` The subnet mask used by the destination host or net.

`Gateway` If the destination is not directly attached to the PowerHub system, this field contains the IP address of the gateway (router) through which packets for the destination are to be routed.

`Met` For entries learned through RIP, this field shows how many hops (routers) away the destination is. For example, if a packet must go through one more router to reach its destination, the metric is `1`.

`State` The state of the route. Possible states are `Up` or `Down`, corresponding to `active` and `inactive`.

`RtSrc` Indicates the source of the routing information:

- `Direct` Indicates that the destination is directly attached to the PowerHub system. Such entries are added automatically when you issue the **`ip add-interface`** command.
- `OExTp1` Indicates that the route was learned through OSPF, from a type-1 External LSA (link-state advertisement). An External LSA indicates that the destination is in another Autonomous System.
- `OExTp2` Indicates that the route was learned through OSPF, from a type-2 External LSA.
- `OInter` Indicates an inter-area route learned through OSPF. The destination is in an area that the PowerHub system is connected to, but the PowerHub system is not in the area that contains the destination. OSPF learns the inter-area routes from summary LSAs.
- `OIntra` Indicates an intra-area route learned through OSPF. The destination is in an area that contains the PowerHub system.
- `Static` Indicates that the route was manually added using the **`ip add-route`** command (see Section 5.11.2 in the *PowerHub Software Manual, V 2.6* (Rev B)).
- `RIP` Indicates that the route was learned through RIP.

`Age` Used only by RIP. Indicates how many seconds have passed since fresh information about this route was received.

`Port` Lists the segment on which packets for this destination should be forwarded. For directly attached nets, a list of segments can appear, because the PowerHub software allows a single net to be used on multiple segments.

3 The OSPF Subsystem

This chapter describes the commands in the **ospf** subsystem. These commands let you display PowerHub OSPF parameters and fine-tune many of those parameters. The OSPF parameters are described briefly in this chapter. For more information about a parameter, see RFC 1583. Where applicable, the descriptions in this chapter refer to specific sections in RFC 1583.

> **NOTE**: Some OSPF parameters can affect routing in the entire autonomous system. Unless you are sure you need to change a specific parameter, FORE Systems recommends that you use the system defaults. (See Section 2.5.1 on page 15.) If you have not yet configured your hub for OSPF, use the procedures in Chapter 2 to do so before continuing with this chapter.

3.1 ACCESSING THE OSPF SUBSYSTEM

To access the **ospf** subsystem, issue the following command at the runtime prompt:

```
ospf
```

3.2 OSPF SUBSYSTEM COMMANDS

Table 3–1 lists and describes the **ospf** subsystem commands and their syntax. For each command, the management capability is listed, as well as the section that contains additional information about the command. (The management capability is explained in the "Getting Started with the User Interface" chapter in the *Installation and Configuration Manual* for your PowerHub system.)

TABLE 3–1 ospf subsystem commands.

Command and Description	Capability*	See…
area|ar add|a *<area-id>* **[***<auth-type>***] [stub-area|sa** *<cost>***]** Adds a new OSPF area.	R	3.3.5
area|ar del|d *<area-id>***|all** Deletes an OSPF area.	R	3.3.5
interface|if *<ip-addr>* **[ar** *<area-id>***] [auth** *<key-str>***] [cost|c** *<cost>***] [priority|p** *<priority>***] [xdelay|x** *<trans-delay>***] [rint|r** *<rxmt-int>***] [hint|h** *<hello-int>***] [rdint|d** *<rtr-dead-int>***] [state|s** *<state>***]** Configures an OSPF interface.	R	3.3.6
net-range|nr add|a *<area-id> <net> <mask>* **[noadv|na]** Adds a network range to an area.	R	3.3.6.1
net-range|nr del|d *<area-id> <net> <mask>* Deletes a network range from an area.	R	3.3.7.1
set|se router-id|ri *<router-id>* Sets the OSPF router ID of the PowerHub system.	R	3.3.1
set|se OSPF-router|ospfr enl|dis Enables or disables OSPF routing.	R	3.3.4
set|se asbd|a enl|dis Enables or disables the PowerHub system as an Autonomous System Border router.	R	3.3.2
set|se auto-vlink|av enl|dis Enables or disables the automatic virtual-link feature.	R	3.3.3
show|sh *<ospf-parm>* Displays OSPF parameters.	R or M	3.4.2

***R = Root; M = Monitor.

TABLE 3–1 ospf subsystem commands.

Command and Description	Capability*	See…
showcfg|scf Displays the OSPF configuration in effect on the PowerHub system.	R or M	3.4.1
stats|s [clear|c] Displays or clears OSPF statistics.	R or M	3.5
virtual-link|vl add|a *<aid> <router-id>* **[auth** *<key-str>***] [xdelay|x** *<trans-delay>***] [rint|r** *<rxmt-int>***] [hint|h** *<hello-int>***] [rdint|d** *<rtr-dead-int>***]** Adds a virtual link.	R	3.3.8
virtual-link|vl del|d *<aid> <router-id>* Deletes a virtual link.	R	3.3.8.1

*R = Root; M = Monitor.

3.3 CONFIGURATION COMMANDS

You can use **ospf** commands to perform the following configuration tasks:

- Specify the PowerHub system's OSPF router ID (**set router-id** command).
- Enable or disable the hub as an Autonomous System Boundary router (**set asbd** command).
- Enable or disable the automatic virtual-link feature (**set auto-vlink** command).
- Enable or disable OSPF routing (**set OSPF-router** command).
- Add an OSPF area (**area add** command).
- Delete an OSPF area (**area del** command).
- Add an OSPF interface to an area (**interface** command).
- Add a network range to an area (**net-range add** command).
- Delete a network range from an area (**net-range del** command).
- Manually add a virtual link (**virtual-link add** command).
- Delete a virtual link (**virtual-link del** command).

The following sections describe the **ospf** configuration commands. Many of these commands are mentioned in the configuration procedures in Section 2.5 on page 14.

3.3.1 Assigning the Hub's OSPF Router ID

To specify the PowerHub system's OSPF router ID, issue the following command.

set|se router-id|ri *<router-id>*

where:

<router-id> — Specifies the ID. Specify the router ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9).

NOTES: You can define the OSPF router ID only when OSPF routing is disabled. To verify that OSPF routing is disabled, issue the **showcfg** command (see Section 3.4.1 on page 43).

FORE Systems recommends that you use the number of one of the IP addresses configured on the PowerHub system as the OSPF router ID. By using one of the IP address on the PowerHub system, you can ensure that OSPF IDs remain unique. Using the number of an IP address as the OSPF router ID does not affect the IP address.

3.3.2 Enabling the Hub as an Autonomous System Border Router

You can enable the PowerHub system to function as an Autonomous System Border router. A hub enabled to be an Autonomous System Border router automatically exports OSPF routes to the networks outside of the OSPF Autonomous System and imports routes from the networks outside the Autonomous System.

NOTE: The PowerHub software exports OSPF routes to RIP and imports RIP routes to OSPF. To control the routes that are exchanged between the protocols, define OSPF import and export filters using the following commands:

- **rip ospf-import-filter**
- **rip ospf-export-filter**

See Section 4.3 on page 65 for information about these commands.

To enable or disable the PowerHub system as an Autonomous System Border router, issue the following command:

set|se asbd|a enl|dis

where:

enl|dis — Specifies whether you want to enable or disable the PowerHub system to function as an Autonomous System Border router. If you specify **enl** (enabled), the PowerHub software can exchange route information between RIP and OSPF. If you specify **dis** (disabled), the software cannot exchange route information. The default is **dis**.

3.3.3 Setting the Automatic Virtual-Link Feature

The automatic virtual-link feature builds virtual links between the areas that are not connected to the backbone. By building the virtual links, the PowerHub software ensures that complete route information reaches all the OSPF routers in the Autonomous System.

> **NOTE**: The automatic virtual-link feature establishes virtual links only between PowerHub systems. To create a virtual link between the PowerHub system and another type of router, use the **`virtual-link add`** command on the PowerHub system. (See Section 3.3.8 on page 41.) See the documentation for your other router for information on establishing that router's end of the virtual link.

The automatic virtual-link feature is enabled by default. To disable (or re-enable) the feature, issue the following command:

`set|se auto-vlink|av enl|dis`

where:

`enl|dis`	Specifies whether you want to enable or disable the automatic virtual-link feature.

3.3.4 Enabling or Disabling OSPF Routing

To enable or disable OSPF routing, issue the following command:

`set|se OSPF-router|ospfr enl|dis`

where:

`enl|dis`	Specifies whether you want to enable or disable OSPF routing. The default is **`dis`** (disabled).

3.3.5 Adding an OSPF Area

When you enable OSPF routing, the PowerHub software automatically creates an OSPF area for the network backbone. The area ID for the backbone is always 0.0.0.0.

Depending upon how you want to organize your network, you might need to add additional OSPF areas. To add an OSPF area to the PowerHub system, issue the following command:

`area|ar add|a` *`<area-id>`* **`[`***`<auth-type>`***`] [stub-area|sa`** *`<cost>`***`]`**

where:

`add|a`	Specifies that you are adding an OSPF area to the PowerHub system.
`<area-id>`	Specifies the area ID. Specify the area ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9). The are ID must be unique within the Autonomous System.

> **NOTE**: The area ID 0.0.0.0 is reserved for the Autonomous System's backbone and is already present in the system.

<auth-type> Specifies the authentication type. You can specify one of the following:

none|no Specifies that the OSPF area you are adding does not use authentication.

simple-password|sp

Specifies that a password is required for OSPF packets sent within this area.

md5|m Specifies that MD5 authentication is required for OSPF packets sent within this area. See RFC 1321 for information about MD5 authentication.

The PowerHub default is **none** (no authentication).

When you add an OSPF interface to this area (using the **interface** command), you specify the actual simple password or MD5 authentication string.

NOTE: All OSPF routers in an area must have the same authentication type and the same authentication string. Also, all OSPF routers on a particular interface should use the same authentication string.

stub-area|sa *<cost>*

Specifies that the area is a stub area. Configuring an area as a stub area reduces OSPF overhead in the network by reducing the amount of OSPF route information flooded to the OSPF routers in the stub area.

The OSPF software does not flood external routing information (information about other Autonomous Systems) into the stub area. Internal routers in the stub area reach Autonomous Systems by using the default route to the stub area's Area Border router.

The OSPF software configures the default route automatically. Note that a stub area's default route is unrelated to the default routes you can define in the **ip** subsystem. OSPF uses the default routes it defines in preference to manually configured default routes.

The cost is the metric for the default route out of the stub area. The stub area's Area Border router advertises the cost as part of the default route. You can specify a value from **1** through **65535**. The default is **1**.

3.3.5.1 Deleting an OSPF Area

To delete an OSPF area, issue the following command:

area|ar del|d *<area-id>*|**all**

where:

del|d	Specifies that you are deleting an OSPF area from the PowerHub system.
<area-id>|**all**	Specifies the area you want to delete. To delete all the OSPF areas defined on this PowerHub system, specify **all**.

> **NOTES**: You must disable OSPF routing before deleting an area. (See Section 3.3.4 on page 35.)
>
> You cannot delete the backbone area (0.0.0.0).

3.3.6 Adding an OSPF Interface to an Area

An OSPF interface is automatically added to the PowerHub system when you add an IP interface. The OSPF interface has the same address as the IP interface. When you enable OSPF routing, the interface is automatically added to the backbone area (0.0.0.0).

In most Autonomous Systems, the PowerHub defaults for the OSPF interface parameters are appropriate for the Autonomous System. However, if you need to change a specific interface parameter, use the following command to do so.

> **NOTE**: You cannot specify the TOS (Type-of-service). The PowerHub software uses TOS 0 (zero, the IP TOS).

interface|it *<ip-addr>* **[ar** *<area-id>***] [auth** *<key-str>***]**
[cost|c *<cost>***] [priority|p** *<priority>***]**
[xdelay|x *<trans-delay>***] [rint|r** *<rxmt-int>***]**
[hint|h *<hello-int>***] [rdint|d** *<rtr-dead-int>***]**
[state|s *<state>***]**

where:

<ip-addr>	Specifies the IP address of the interface. Specify the interface in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9). The IP address must already be present in the hub's IP interface table before you can use it to create an OSPF interface. To display a list of all the IP interfaces configured on the PowerHub system, issue the **ip route-table** command. (See Section 2.7.3 on page 29.) To add an IP interface to the PowerHub system, use the **ip add-interface** command. (See Section 2.5.3 on page 16.)

ar *<area-id>*

Specifies the OSPF area in which you are placing the OSPF interface. An OSPF interface can belong to only one area. The area must already be configured (using the **add area** command; see Section 3.3.5 on page 35).

auth *<key-str>*

Specifies the authentication string. For a simple password, you can specify any combination of up to eight numbers, letters, and special characters. For MD5 authentication, you can specify any combination of up to 16 numbers, letters, and special characters. The authentication string is case-sensitive.

If the area to which you are adding this interface does not require an authentication string, use empty quotation marks (**""**).

cost|c *<cost>*

Specifies the cost of using this interface. The PowerHub software advertises the cost in Router Links Advertisements. You can specify a cost from **1** through **32**. This parameter does not have a default value. The cost depends upon the wire speed of the segment on which you are adding this interface. Unless you are certain that you need to change the cost, FORE Systems recommends that you omit this argument and use the value determined by the PowerHub software.

priority|p *<priority>*

Specifies this interface's priority during the election process for the DR (Designated Router). The interface with the highest priority number is elected as the DR. The interface with the second-highest priority number is elected as the BDR (Backup Designated Router).

You can specify a priority from **0** through **255**. Priority increases from **1** (lowest) to **255** (highest). A priority of **0** (zero) makes this interface ineligible for becoming the DR. The default is **1**.

If all OSPF interfaces within an Autonomous System have the same priority, the DR and BDR are elected based on the interface addresses. The interface with the highest OSPF address is elected as the DR. The interface with the second-highest OSPF address is elected as the BDR.

NOTE: Generally, an OSPF router has only one interface per area. If the PowerHub system has multiple interfaces to the same area, the interface priority still applies.

xdelay|x *<trans-delay>*

Specifies the interface *transmission delay*, which is the estimated number of seconds it takes to transmit a Link State Update packet over this interface. The PowerHub software adds the transmission delay you specify to the ages of the LSAs contained in the Link State Update packets sent on this interface.

You specify a delay from **1** through **3600**. The default is **1**.

See Section 9 in RFC 1583 for information about choosing the transmission delay.

rint|r *<rxmt-int>* Specifies the retransmission interval. The *retransmission interval* is the number of seconds between transmissions of LSAs to the OSPF routers adjacent to this interface. The retransmission interval also is used when transmitting Database Description and Link State Request packets.

You can specify an interval from **1** through **3600**. The default is **5**.

hint|h *<hello-int>* Specifies the hello interval. The *hello interval* is the number of seconds between transmission of Hello packets on this interface. You can specify an interval from**1** through **65536**. The default is **10**.

NOTE: The hello interval (**hint**) and the router-dead interval (**rdint**) must match on neighbors. That is, the PowerHub values for these parameters must match the values on the hub's neighbor for these parameters. If the hub's OSPF neighbor also is a PowerHub system, you can ensure that the values match by accepting the defaults for these parameters. If the neighbor is not a PowerHub system, you might need to change the value on the neighbor or on the hub so that the values on both routers match.

rdint|d *<rtr-dead-int>*

Specifies the router-dead interval. The *router-dead interval* is the number of seconds the hub's OSPF neighbors should wait before declaring that the hub (as an OSPF router) is down.

You can specify a router-dead interval from **1** through **65536**. Specify an interval that is an even multiple of the Hello interval. The default is **40**.

state|s *<state>* Specifies the state of the OSPF interface. You can specify **enable** or **disable**. The default is **enable**.

NOTE: The OSPF interfaces are automatically enabled by the software when you enable OSPF routing.

3.3.6.1 Removing an Interface from an Area

To remove an interface from an OSPF area, disable the interface using the following form of the **interface** command:

interface|if *<ip-addr>* **state|s disable**

This command does not remove the interface from the PowerHub system. The interface still appears in the IP interface table. (To display the IP interface table, issue the **ip interface-table** command.)

3.3.7 Adding a Network Range to an Area

You do not need to add network ranges to OSPF areas. The PowerHub software automatically advertises all the networks on all the OSPF interfaces on the hub to other OSPF routers.

You can add network ranges to reduce OSPF overhead or to hide certain networks from other OSPF routers.

When you add a network range to an area, link-state information for the networks within the range is summarized in the LSAs sent by the hub to its OSPF neighbors. Therefore, if you have many networks within an area, adding the networks as a network range can help reduce OSPF overhead.

In addition, you can use the **noadv** argument with the **net-range** command to prevent the hub from advertising routes to the networks within a network range. When the hub sends LSAs to its neighbors, LSAs for the networks in the hidden network range are not sent to the hub's neighbors. Therefore, other routers in the Autonomous System do not learn about the hidden networks.

> **NOTE**: None of the networks within the network range you add to an area can be in other areas.

To add a network range to an OSPF area, issue the following command:

net-range|nr add|a *<area-id>* *<net>* *<mask>* **[noadv|na]**

where:

<area-id>	Specifies the OSPF area. The area must already be added to the system. To add an area, use the **add area** command. (See Section 3.3.5 on page 35.)
<net>	Specifies an IP network. Specify the network address in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9). The address you specify is ANDed with the subnet mask you specify for the *<mask>* argument.
<mask>	Specifies the IP mask associated with the IP network address you specified for the *<net>* argument. The mask indicates the portion of the IP network address that is to be regarded as the network portion of the address. Specify the mask in dotted decimal notation (ex: 255.255.255.0).
noadv|na	Prohibits the OSPF software from advertising this network range in the LSAs transmitted by the hub to its OSPF neighbors. If you use this argument, other OSPF routers do not learn about the presence of the network range.

Here is an example of this command. In this example, the network range specified by IP address 200.200.200.200 and subnet mask 255.255.255.0 is added to area 1.1.1.1. When area 1.1.1.1 sends LSAs to other areas, the LSAs will contain summary information for the networks within the network range, instead of detailed link-state information for each network within the network range.

```
9:PowerHub:ospf# net-range add 1.1.1.1 200.200.200.0 255.255.255.0
OSPF: Net "200.200.200.0" with Mask "255.255.255.0" added to area
"1.1.1.1"
```

If the **noadv** argument had been specified with the command, the area would not report the networks within the specified network range.

3.3.7.1 Deleting a Network Range from an Area

To delete a network range, issue the following command:

net-range|nr del|d *<area-id> <net> <mask>*

where:

<area-id>	Specifies the OSPF area.
<net>	Specifies the IP network address.
<mask>	Specifies the subnet mask associated with the IP address.

Here is an example of this command.

```
10:PowerHub:ospf# net-range del 1.1.1.1 200.200.200.0 255.255.255.0
OSPF: Net "200.200.200.0" with Mask "255.255.255.0" deleted from area
"1.1.1.1"
```

After you delete a network range, the PowerHub system sends detailed link-state information for each network, instead of summarizing the link-state information for the entire range.

3.3.8 Adding a Virtual Link

Depending upon how you configure your OSPF network, it is possible for some areas to be completely disconnected from one another. Areas become disconnected from one another when they are not attached to the backbone and do not share a Border router.

The PowerHub software can automatically link disconnected areas using the automatic virtual-link feature. This feature links together PowerHub systems configured as OSPF routers when those systems are separated from one another. (See Section 3.3.3 on page 35.)

If some of the OSPF routers in your Autonomous System are not PowerHub systems, you can link areas that are separated by defining a virtual link between the areas. The virtual link makes the disconnected areas virtual neighbors. LSAs from an area reach that area's virtual neighbor by travelling through a transit area. The transit

area is an area between the two virtual neighbors that passes traffic between the neighbors.

> **NOTE**: You must add the transit area to your OSPF network before configuring the virtual link.

To add a virtual link, use the following command:

virtual-link|vl add|a *<aid> <router-id>* **[auth** *<key-str>***]**
[xdelay|x *<trans-delay>***]**
[rint|r *<rxmt-int>***]**
[hint|h *<hello-int>***]**
[rdint|d *<rtr-dead-int>***]**

The values and defaults for these arguments are the same as the arguments and defaults for the **interface** command. (See Section 3.3.6 on page 37.)

3.3.8.1 Deleting a Virtual Link

To delete a virtual link, issue the following command:

virtual-link|vl del|d *<aid> <router-id>*

where:

<aid>	Specifies the area ID of the transit area. Specify the area ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9).
<router-id>	Specifies the OSPF Router ID of the virtual neighbor. Specify the router ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9).

> **NOTE**: You can use the **virtual-link del** command to delete a virtual link created by the software automatically using the automatic virtual-link feature. However, if the automatic virtual-link feature is enabled, the software adds the link again. To prevent the software from adding a virtual link again, disable the automatic virtual-link feature by issuing the **set auto-vlink dis** command. (See Section 3.3.3 on page 35.)

3.4 CONFIGURATION-INFORMATION COMMANDS

After you configure your PowerHub system for OSPF routing, you can verify the configuration using the following commands:

showcfg	Displays basic information such as whether memory has been allocated for OSPF and whether OSPF forwarding is enabled.
show *<ospf-parm>*	Displays detailed information about specific OSPF parameters such as areas and interfaces.

3.4.1 Displaying Basic Configuration Settings

To display the hub's OSPF configuration, issue the following command:

showcfg|scf

This command indicates whether memory has been allocated for OSPF and indicates the state (enabled or disabled) of OSPF features, as shown in the following example.

```
11:PowerHub:ospf# showcfg
OSPF Router                          : Memory Available
OSPF Routing                         : Enabled
OSPF Router ID                       : 1.1.1.1
OSPF Version Number                  : 2
OSPF Autonomous System Border Router : Enabled
Automatic Virtual Link Feature       : Enabled
```

The fields in this display show the following information:

OSPF Router	Indicates whether you have allocated a portion of the PowerHub memory for OSPF. In this example, memory has been allocated. If memory has not been allocated, you can allocate the memory using the following command: **main getmem ospf**. (See Section 2.5.2 on page 15.)
OSPF Routing	Indicates whether OSPF routing is enabled. If OSPF routing is disabled, you can enable it using the **set OSPF-routing enl** command. (See Section 3.3.4 on page 35.)
OSPF Router ID	Displays the OSPF router ID you assigned to this PowerHub system. If no router ID is displayed, you can assign the router ID using the **set router-id** command. (See Section 3.3.1 on page 34.)
OSPF Version Number	Indicates that OSPF version 2 is implemented on the PowerHub system. This value does not change.

`OSPF Autonomous System Border Router`

Indicates whether you have configured the PowerHub system as an Autonomous System Border router. This feature is disabled by default. To enable the feature, issue the **`set asbd enl`** command. (See Section 3.3.2 on page 34.)

`Automatic Virtual Link Feature`

Indicates whether the automatic virtual-link feature is enabled. This feature is enabled by default. To disable the feature, issue the **`set auto-link dis`** command. (See Section 3.3.3 on page 35.)

3.4.2 Displaying Specific Configuration Parameters

To display specific OSPF configuration parameters, issue the following command:

`show|sh` *`<ospf-parm>`*

where:

`<ospf-parm>` Specifies one of the following:

`area|ar [`*`<aid>`***`]`**

Displays information about the OSPF areas to which this PowerHub system belongs. If you specify an area ID (*`<aid>`*), information for only that area is displayed.

`interface|if [`*`<ip-addr>`***`]`**

Displays information for the OSPF interfaces configured on this PowerHub system. If you specify an interface address (*`<ip-addr>`*), information for only that interface is displayed.

`lsdb|l [`*`<lsdbid> <rid> <type> <aid>`***`]`**

Displays the contents of LSAs (Link State Advertisements). You can use the following arguments to display detailed information about an advertisement:

`<lsdbid>` A Link State Database Identifier in dotted-decimal notation (xxx.xxx.xxx.xxx, where each “x” is an integer from 0 – 9).

`<rid>` An originating Autonomous System Router ID in dotted-decimal notation (xxx.xxx.xxx.xxx, where each “x” is an integer from 0 – 9).

<type> An LSA type. You can specify one of the following:

r Router LSA

n Network LSA

s Summary LSA

a Autonomous System Summary LSA

e External LSA

<aid> Area ID, in dotted-decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9).

neighbor|nb

Displays information about this PowerHub system's OSPF neighbors. Only directly-attached neighbors are listed. Virtual neighbors are not listed. A virtual neighbor is an OSPF router that is reached using a virtual link.

net-range|nr [*<aid>***]**

Displays the network ranges defined for an area.

virtual-link|vl [*<aid> <router-id>***]**

Displays information about Virtual Links. You can use the optional arguments to display detailed information about the virtual link for a specific neighbor.

The following sections show examples of the information you can display using the **show** command.

3.4.2.1 Area Information

Use the **show area [***<aid>***]** command to display information about the OSPF areas configured on the PowerHub system.

Here are some examples of the information displayed by this command. In the following example, information is displayed for all the OSPF areas configured on the PowerHub system.

```
12:PowerHub:ospf# show area
                Auth Import AS Number of  # Area # AS   Number of  Stub Area
Area Id         Type Ext LSAs  Spf Runs   Bdr    Bdr    Area LSAs  Cost
--------------- ---- --------- ---------- ------ ------ ---------- --------
0.0.0.0          no  Enabled   12         4      4      13         -----
1.1.1.1          no  Enabled   12         2      2      15         -----
1.2.3.4          md5 Enabled   12         0      0      0          -----
2.3.4.5          sp  Enabled   12         0      0      0          -----
3.3.3.3          no  Enabled   12         1      1      16         -----
33.0.33.0        no  Enabled   12         0      0      0          -----
33.33.33.33      no  Enabled   12         0      0      0          -----
```

In the following example, information is displayed for a specific OSPF area.

```
13:PowerHub:ospf# show area 1.2.3.4
                Auth Import AS Number of  # Area # AS   Number of  Stub Area
Area Id         Type Ext LSAs  Spf Runs   Bdr    Bdr    Area LSAs  Cost
--------------- ---- --------- ---------- ------ ------ ---------- --------
1.2.3.4          md5 Enabled   12         0      0      0          -----
```

The fields in this display show the following information:

Area ID — Displays the OSPF area ID you assigned using the **area add** command. The area ID is a 32-bit integer expressed in dotted decimal notation. The area ID 0.0.0.0 is the backbone area ID and is added automatically by the PowerHub software.

Auth Type — Displays the authentication type you assigned for this area using the **area add** command. The authentication type can have one of the following values:

- `no` No authentication is required for this area.
- `sp` A simple password is required for this area.
- `md5` MD5 authentication is required in this area. See RFC 1321 for information about MD5.

Import AS Ext LSAs — Specifies whether this area is configured to import external LSAs from other Autonomous Systems. The value can be `Enabled` or `Disabled`. To change the state of this parameter, use the **set asbd enl|dis** command. (See Section 3.3.2 on page 34.)

Number of SPF Runs — Indicates the number of times the PowerHub software has calculated this area's intra-area route table. This number is reset to zero if you disable OSPF routing, reboot the software, or power down the PowerHub system.

 # Area Bdr — Indicates the number of Area Border routers that can be reached from this area.

 # AS Bdr — Indicates the number of Autonomous System Border routers that can be reached from this area.

Number of Area LSAs — Indicates the number of LSAs in this area's LSA database. This number does not include external LSAs.

Stub Area Cost — If this area is a stub area, the metric for the stub area is indicated in this field. If this area is not a stub area, this field contains dashes (`-----`). You assign a stub area's metric when you add the area using the **area add** command.

3.4.2.2 Interface Information

Use the **show interface [***<ip-addr>***]** command to display information about an OSPF interface. If you omit the optional *<ip-addr>* argument, summary information is displayed for all OSPF interfaces configured on this PowerHub system. If you use the *<ip-addr>* argument, detailed information about the OSPF address you specify is displayed.

Here are some examples of the information displayed by this command. In the first example, summary information for all OSPF interfaces on the hub is displayed.

```
14:PowerHub:ospf# show interface
IP Address      Area Id         DR              BDR             Admin
--------------- --------------- --------------- --------------- ---------
192.9.222.1     0.0.0.0         0.0.0.0         0.0.0.0         Enabled
150.1.200.1     0.0.0.0         0.0.0.0         0.0.0.0         Enabled
129.213.72.1    0.0.0.0         129.213.72.2    129.213.72.1    Enabled
150.1.100.1     1.1.1.1         150.1.100.3     150.1.100.1     Enabled
111.222.111.1   3.3.3.3         0.0.0.0         0.0.0.0         Enabled
```

The fields in this display show the following information:

`IP Address`	The IP address of the OSPF interface.
`Area ID`	The OSPF area to which the interface belongs. The area ID 0.0.0.0 indicates the backbone.
`DR`	The OSPF router ID of the DR (Designated Router).
`BDR`	The OSPF router ID of the BDR (Backup Designated Router). The BDR is elected when OSPF is enabled.
`Admin`	Indicates whether the interface is enabled or disabled. The PowerHub software automatically adds an OSPF interface when you add an IP interface (using the **ip add-interface** command). To enable or disable an OSPF interface, use the **state** argument with the **interface** command. (See Section 3.3.6 on page 37.)

In the following example, detailed information is displayed for a specific OSPF interface.

```
15:hub2:ospf# show interface 192.9.222.1
IP Address                  : 192.9.222.1
Area ID                     : 0.0.0.0
Interface Type              : Broadcast
Administrative Status       : Enabled
Router Priority             : 1
Interface Metric            : 10
Transit Delay               : 1
Retransmission Interval     : 5
Hello Interval              : 10
Router Dead Interval        : 40
Interface State             : Down
Designated Router           : 0.0.0.0
Backup Designated Router    : 0.0.0.0
Interface Events            : 0
Authorization Key String    : MKT
Authorization Failures      : 0
```

The fields in this display show the following information:

`IP Address` The IP address of the OSPF interface.

`Area ID` The OSPF area to which the interface belongs. (The area ID 0.0.0.0 indicates the backbone.)

`Interface Type` The type of OSPF interface. The interface can be one of the following types:

- `Broadcast` The network to which the interface is attached is a broadcast network. That is, the network supports more than two attached routers and those routers have the ability to address a single physical message (broadcast) to all of the attached routers.
- `Virtual Link` The network spans multiple non-backbone areas.

`Administrative Status`

The state of the OSPF interface. The state can be one of the following:

- `Enabled` The interface can participate in OSPF routing.
- `Disabled` The interface cannot participate in OSPF routing.

The PowerHub software automatically adds an OSPF interface when you add an IP interface (using the **ip add-interface** command). To enable or disable an OSPF interface, use the state argument with the **interface** command. (See Section 3.3.6 on page 37.)

`Router Priority` The priority used by the PowerHub software when electing the DR and the BDR. A priority of 0 (zero) makes this interface ineligible for becoming the DR or the BDR.

You can change the priority using the *<priority>* argument of the **interface** command. (See Section 3.3.6 on page 37.)

`Interface Metric` The cost of using this interface. You can change the cost using the *<cost>* argument with the **interface** command. (See Section 3.3.6 on page 37.)

`Transit Delay` The interface transmission delay for this interface. You can change the transmission delay using the *<trans-delay>* argument with the **interface** command. (See Section 3.3.6 on page 37.)

`Retransmission Interval`

The retransmission interval for this interface. You can change this parameter using the *<rxmt-int>* argument with the **interface** command. (See Section 3.3.6 on page 37.)

`Hello Interval` The Hello interval for this interface. You can change this parameter using the *<hello-int>* argument with the **interface** command. (See Section 3.3.6 on page 37.)

`Router Dead Interval`

The router-dead interval for this interface. You can change this parameter using the *<rtr-dead-int>* argument with the **interface** command. (See Section 3.3.6 on page 37.)

`Interface State` The forwarding state of the PowerHub interface on which this OSPF interface is configured. The state can be one of the following:

- `Down` The segment cable is faulty or has been unplugged, or the segment has been disabled by the **`bridge port`** command or by the automatic segment-state detection feature.
- `Up` The link is good and the segment is enabled.

`Designated Router` The OSPF router ID of the DR.

`Backup Designated Router`

The OSPF router ID of the BDR.

`Interface Events` The number of events that have caused a change in this interface's state. See Section 9.2 in RFC 1583.

`Authorization Key String`

The authorization string (password) for the interface. The authorization string is specified by the *`<key-str>`* argument of the **`interface`** command. If this field is blank, then no authorization string is required for this interface. (See Section 3.3.5 on page 35.)

`Authorization Failures`

The number of packets that could not be forwarded because they failed authorization verification. A packet can fail authorization verification when the packet's authorization string does not match the authorization string you specified for an area. (See Section 3.3.5 on page 35.)

3.4.2.3 Link-State Database Information

Use the following command to display information about a link-state database:

`show lsdb|l [`*`<lsdbid> <rid> <type> <aid>`***`]`**

where:

`<lsdbid>` Specifies the ID of a specific LSA (link-state advertisement).

`<rid>` Specifies the OSPF router ID of the router from which the link-state database was received.

<type> Specifies the LSA type, which can be one of the following types:

- **r** Router LSA
- **n** Network LSA
- **s** Summary LSA
- **a** Autonomous System Summary LSA
- **e** External LSA

<aid> Specifies the ID of the area to which the LSA applies.

If you omit the optional arguments, summary information is displayed for all the LSAs present in the PowerHub system's LSA database. To display detailed information about a specific LSA, use the optional arguments.

Here are some examples of the information displayed by this command. In the first example, summary information for all LSAs in the hub's LSA database is displayed.

```
16:PowerHub:ospf# show lsdb
Area Id         Lsdb Type      Link State ID   Router ID       Sequence
--------------- -------------- --------------- --------------- -----------
0.0.0.0         routerLink     1.1.1.1         1.1.1.1         -2147483552
0.0.0.0         routerLink     2.2.2.2         2.2.2.2         -2147483303
0.0.0.0         routerLink     3.3.3.3         3.3.3.3         -2147483615
0.0.0.0         routerLink     5.5.5.5         5.5.5.5         -2147483576
0.0.0.0         networkLink    80.100.1.3      3.3.3.3         -2147483635
0.0.0.0         networkLink    129.213.72.2    5.5.5.5         -2147483635
0.0.0.0         summaryLink    87.0.0.0        2.2.2.2         -2147483348
0.0.0.0         summaryLink    150.1.100.0     1.1.1.1         -2147483578
0.0.0.0         summaryLink    150.1.100.0     3.3.3.3         -2147483632
0.0.0.0         summaryLink    152.16.0.0      3.3.3.3         -2147483631
0.0.0.0         summaryLink    170.170.1.0     3.3.3.3         -2147483630
0.0.0.0         asSummaryLink  1.1.1.1         3.3.3.3         -2147483635
0.0.0.0         asSummaryLink  3.3.3.3         1.1.1.1         -2147483634
1.1.1.1         routerLink     1.1.1.1         1.1.1.1         -2147483638
1.1.1.1         routerLink     3.3.3.3         3.3.3.3         -2147483635
1.1.1.1         networkLink    150.1.100.3     3.3.3.3         -2147483646
1.1.1.1         summaryLink    44.0.0.0        3.3.3.3         -2147483640
1.1.1.1         summaryLink    80.100.0.0      3.3.3.3         -2147483640
1.1.1.1         summaryLink    80.200.0.0      3.3.3.3         -2147483640
1.1.1.1         summaryLink    87.0.0.0        3.3.3.3         -2147483640
1.1.1.1         summaryLink    129.213.0.0     1.1.1.1         -2147483642
1.1.1.1         summaryLink    152.16.0.0      3.3.3.3         -2147483640
1.1.1.1         summaryLink    170.170.1.0     3.3.3.3         -2147483640
1.1.1.1         summaryLink    180.180.0.0     3.3.3.3         -2147483640
1.1.1.1         summaryLink    200.200.200.0   3.3.3.3         -2147483640
1.1.1.1         summaryLink    213.100.45.0    3.3.3.3         -2147483640
1.1.1.1         asSummaryLink  2.2.2.2         3.3.3.3         -2147483640
1.1.1.1         asSummaryLink  5.5.5.5         1.1.1.1         -2147483645
<example truncated here for brevity>
```

The fields in this display show the following information:

`Area ID` The OSPF area from which the LSA was received.

`Lsdb Type` The type of LSA.

`Link State ID` The ID of the LSA, in dotted-decimal notation. The LSA ID is determined by the type of the LSA:

LSA Type	LSA ID
An Internal router's LSA (`routerLink`).	The originating router's OSPF router ID.
A network LSA (`networkLink`).	The IP interface address of the network's DR (Designated Router).
A summary LSA (`summaryLink`).	The destination network's IP address.
An Autonomous System Border router's LSA (`asSummaryLink`).	The OSPF router ID of the Autonomous System Boundary router described by the LSA.
An Autonomous System Border router's external LSA (`asExternalLink`).	The destination network's IP address.

`Router ID` The OSPF router from which the LSA was received.

`Sequence` The sequence number of the LSA. The sequence number is a 32-bit signed integer. A higher sequence number indicates a more recent LSA. You can use the LSA sequence numbers to detect old or duplicate LSAs.

In the following example, detailed information is displayed about a specific LSA.

```
17:PowerHub:ospf# show lsdb 1.1.1.1 1.1.1.1 r 0.0.0.0
Detailed View
Area ID                     : 0.0.0.0
Link State Database Type    : routerLink
Link State ID               : 1.1.1.1
Originating Router ID       : 1.1.1.1
Sequence Number             : -2147483552
Advertisement Age           : 1503
Advertisement Checksum      : ccac
The OSPF Link State Database Advertisement: (26 per line)
00 00 02 01 01 01 01 01 01 01 01 01 80 00 00 60 cc ac 00 30 03 00 00 02 81 d5
48 02 81 d5 48 01 02 00 00 0a 03 03 03 03 96 01 64 01 04 00 00 0a
```

The fields in this display show the following information:

`Area ID` The OSPF area from which the LSA was received.

`Link State Database Type`

The type of LSA. The LSA can be one of the following types:

`routerLink`	Internal router LSA.
`networkLink`	Network LSA.
`summaryLink`	Summary LSA.
`asSummaryLink`	Autonomous System Border router LSA.
`asExternalLink`	External LSA.

`Link State ID` The ID of the LSA. The LSA ID depends upon the type of the LSA:

LSA Type	LSA ID
An internal router's LSA (`routerLink`).	The originating router's OSPF router ID.
A network LSA (`networkLink`).	The IP interface address of the network's DR (Designated Router).
A summary LSA (`summaryLink`).	The destination network's IP address.
An Autonomous System Border router's LSA (`asSummaryLink`).	The OSPF router ID of the Autonomous System Border router described by the LSA.
An Autonomous System Border router's external LSA (`asExternalLink`).	The destination network's IP address.

`Originating Router ID`

The OSPF router from which the LSA was received.

`Sequence Number` The sequence number of the LSA. The sequence number is a 32-bit signed integer. A higher sequence number indicates a more recent LSA. You can use the LSA sequence numbers to detect old or duplicate LSAs.

`Advertisement Age` The age, in seconds, of the LSA.

`Advertisement Checksum`

The checksum for the LSA.

`The OSPF Link State Database Advertisement`

The contents of the LSA, in hexadecimal digits.

3.4.2.4 Neighbor Information

Use the **`show neighbor`** command to display information about the PowerHub system's OSPF neighbors. Here is an example of the information displayed by this command.

```
18:PowerHub:ospf# show neighbor
IP Address      Router ID       Pri State    Events RTrQ
--------------- --------------- --- -------- ------ -----
129.213.72.2    5.5.5.5         1   full     6      0
150.1.100.3     3.3.3.3         1   full     6      0
```

The fields in this display show the following information:

`IP Address` — The IP address of the neighbor's interface to the PowerHub system.

`Router ID` — The ID of the OSPF router that contains the neighbor.

`Pri` — The priority of the OSPF router that contains the neighboring interface. The priority is used when the PowerHub software elects a Designated router and a Backup Designated router. If the priority is 0 (zero), the OSPF router is ineligible to become the Designated router or Backup Designated router.

`State` — The state of the relationship with the neighboring interface's router. The state can be one of the following:

- `down` — The hub has not received recent information from the neighbor.
- `attempt` — The hub has not received recent information from the neighbor, but the software is attempting to contact the neighbor by sending Hello packets. You can change the Hello interval using the **`hint`** argument of the **`interface`** command. (See Section 3.3.6 on page 37.)
- `init` — The hub recently received a Hello packet from the neighbor.
- `two Way` — Communication between the hub and the neighbor now is bi-directional.
- `ex start` — The hub and its neighbor are beginning to exchange their link-state databases.

`exchange`	The hub is sending its link-state database to the neighbor.
`loading`	The hub is sending Link State Request packets to the neighbor, requesting the LSAs that are more recent than the information contained in the link-state database the hub sent to that neighbor. The hub updates its link-state database with the new LSAs received from the neighbor.
`full`	The hub and the neighbor have finished exchanging their link-state databases. For more information about these states, see Section 10.1 in RFC 1583.

`Events` The number of times the state of the neighbor relationship (see the `State` field) has changed. See Section 10.1 in RFC 1583.

`RTrQ` The current length of the retransmission queue.

3.4.2.5 Network-Range Information

Use the **`show net-range [`***`<aid>`***`]`** command to display information about the network ranges assigned to the areas configured on the PowerHub system. If you omit the optional ***`<aid>`*** argument, summary information is displayed for all the network ranges in all the areas on the hub. To display network-range information for a specific area, use the ***`<aid>`*** argument.

Here is an example of the information displayed by the **`show net-range`** command. In this example, the optional ***`<aid>`*** argument is omitted. Only one network range is listed in the display, indicating that only one OSPF network range has been configured on the hub.

```
19:PowerHub:ospf# show net-range
Area ID          Net              Mask             Advertise
--------------   --------------   --------------   ---------
1.1.1.1          200.200.200.0    255.255.255.0    Enabled
```

The fields in this display show the following information:

`Area ID` The OSPF area that contains the network range.

`Net` The IP address of the network or subnet portion of the network range. The network number is ANDed with the subnet mask (see the `Mask` field) to make the network range.

`Mask` The subnet mask that is ANDed with the network number (see the `Net` field) to make the network range.

`Advertise` Indicates whether this network range is advertised to other areas. The advertise state can be `Enabled` or `Disabled`. The advertise state is enabled by default. To prevent the hub from advertising the network range to other areas, use the **`noadv`** argument with the **`net-range`** command. (See Section 3.3.7 on page 40.)

3.4.2.6 Virtual-Link Information

Use the following command to display information about a virtual link:

`show virtual-link|vl [`*`<aid> <router-id>`***`]`**

If you omit the optional arguments, summary information is displayed for all the virtual links that exist between this PowerHub system and other OSPF routers. To display detailed information about a virtual link, use the optional arguments.

Here are some examples of the information displayed by this command. In the first example, summary information is displayed. The hub in this example has only one virtual link to another OSPF router.

```
20:PowerHub:ospf# show virtual-link
Area ID         Router ID       IP Address      If State Nbr State
--------------- --------------- --------------- -------- ---------
1.1.1.1         3.3.3.3         150.1.100.3     up       full
```

The fields in this display show the following information:

`Area ID` The OSPF area on the local side of the virtual link.

`Router ID` The router ID of the OSPF router on the local end of the virtual link. (The PowerHub system's OSPF router ID.)

`IP Address` The IP address of the router on the remote end of the Virtual Link. Routers can have many IP addresses. This IP address is the one assigned to the remote router's segment that connects the remote router to the PowerHub system.

`IF State` The state of the virtual interface. The state can be one of the following:

- `up` The interface can be used to send and receive OSPF route information.
- `down` The interface is unavailable for sending or receiving OSPF traffic. The interface's link state will be reported as down in LSAs sent from this OSPF router.

`Nbr State` The state of the relationship between the PowerHub system and the OSPF router at the remote end of the virtual interface. The state can be one of the following:

`down`	The hub has not received recent information from the neighbor.
`attempt`	The hub has not received recent information from the neighbor, but the software is attempting to contact the neighbor by sending Hello packets. You can change the Hello interval using the **`hint`** argument of the **`interface`** command. (See Section 3.3.6 on page 37.)
`init`	The hub recently received a Hello packet from the neighbor.
`two Way`	Communication between the hub and the neighbor now is bi-directional.
`ex start`	The hub and its neighbor are beginning to exchange their link-state databases.
`exchange`	The hub is sending its link-state database to the neighbor.
`loading`	The hub is sending Link State Request packets to the neighbor, requesting the LSAs that are more recent than the information contained in the link-state database the hub sent to that neighbor. The hub updates its link-state database with the new LSAs received from the neighbor.
`full`	The hub and the neighbor have finished exchanging their link-state databases.

For more information about these states, see Section 10.1 in RFC 1583.

In the following example, detailed information is displayed for a specific virtual link.

```
21:PowerHub:ospf# show virtual-link 1.1.1.1 3.3.3.3
Area ID                           : 1.1.1.1
Router ID                         : 3.3.3.3
IP Address                        : 150.1.100.3
Transit Delay                     : 1
Retransmission Interval           : 5
Hello Interval                    : 10
Router Dead Interval              : 60
Authorization Key String          :
Authorization Failures            : 0
Virtual Interface State           : up
Virtual Interface Events          : 1
Virtual Neighbor State            : full
Virtual Neighbor Events           : 5
Virtual Neighbor Retransmission Que : 0
```

The fields in this display show the following information:

`Area ID` — The OSPF area on the local side of the virtual link.

`Router ID` — The router ID of the OSPF router on the local end of the virtual link. (The PowerHub system's OSPF router ID.)

`IP Address` — The IP address of the router on the remote end of the Virtual Link. Routers can have many IP addresses. This IP address is the one assigned to the remote router's segment that connects the remote router to the PowerHub system.

`Transit Delay` — The interface transmission delay for this interface.

`Retransmission Interval` — The retransmission interval for this PowerHub interface.

`Hello Interval` — The Hello interval for this PowerHub interface.

`Router Dead Interval` — The router dead interval for this PowerHub interface.

`Authorization Key String` — The authorization string for the PowerHub interface. The authorization string is specified by the *<key-str>* argument of the **`interface`** command. If this field is blank, then no authorization string is required for this interface.

`Authorization Failures` — The number of times another OSPF router tried to use this interface but did not supply the correct authorization string.

`Virtual Interface State`

The state of the virtual interface. The state can be one of the following:

`up`	The interface can be used to send and receive OSPF route information.
`down`	The interface is unavailable for sending or receiving OSPF route information. The interface's link state will be reported as down in LSAs sent from this OSPF router.

`Virtual Interface Events`

The number of times the state (see the `Virtual Interface State` field) has changed since OSPF routing was enabled.

`Virtual Neighbor State`

The state of the relationship with the OSPF router on the remote end of the virtual link. The state can be one of the following:

`down`	The hub has not received recent information from the neighbor.
`attempt`	The hub has not received recent information from the neighbor, but the software is attempting to contact the neighbor by sending Hello packets. You can change the Hello interval using the **`hint`** argument of the **`interface`** command. (See Section 3.3.6 on page 37.)
`init`	The hub recently received a Hello packet from the neighbor.
`two Way`	Communication between the hub and the neighbor now is bi-directional.
`ex start`	The hub and its neighbor are beginning to exchange their link-state databases.
`exchange`	The hub is sending its link-state database to the neighbor.
`loading`	The hub is sending Link State Request packets to the neighbor, requesting the LSAs that are more recent than the information contained in the link-state database the hub sent to that neighbor. The hub updates its link-state database with the new LSAs received from the neighbor.

`full` The hub and the neighbor have finished exchanging their link-state databases.

For more information about these states, see Section 10.1 in RFC 1583.

`Virtual Neighbor Events`

The number of times the relationship with the remote end of the virtual link has changed since OSPF routing was enabled. The state is displayed in the `Virtual Neighbor State` field.

`Virtual Neighbor Retransmission Que`

The current length of the retransmission queue on the OSPF router at the remote end of the virtual link.

3.5 STATISTICS COMMAND

As soon as you enable OSPF forwarding, the PowerHub software begins collecting OSPF statistics. Use the **`stats [clear]`** command to display or clear the statistics.

3.5.1 Displaying OSPF Statistics

To display the OSPF statistics, issue the following command:

`stats|s`

Here is an example of the information displayed by the **`stats`** command.

```
22:PowerHub:ospf# stats
External Link-State Advertisements        : 0
Checksum of the External LSA Database     : 0
New Link-State Advertisements originated  : 105
Link-State Advertisements received        : 121
OSPF Area Border Router                   : True
Total Authorization Failures              : 0
ospf: Total 100 KB. Used 33 KB. Free  66 %
```

The fields in this display show the following information:

`External Link-State Advertisements`

The total number of type-1 and type-2 External LSAs in the PowerHub system's OSPF link-state database.

`Checksum of the External LSA Database`

A summation of the checksums of all the type-1 and type-2 External LSAs in the PowerHub system's link-state database.

`New Link-State Advertisements originated`

The number of LSAs originated by the PowerHub system.

`Link-State Advertisements received`

The number of LSAs received by the PowerHub system.

`OSPF Area Border Router`

Indicates whether the PowerHub system is functioning as an Area Border router. This field can have the following values:

`True` Indicates that two or more areas are active in the PowerHub system.

`False` Indicates that only one area is active in the PowerHub system. The PowerHub system is not functioning as an Area Border router.

`Total Authorization Failures`

Memory allocation failures. If this field contains a value higher than "0," you need to allocate more memory to the **ospf** subsystem. (See the `ospf:` field and Section 2.5.2 on page 15.) Note that this field is not related to the authorization key.

`ospf:` The amount of memory currently in use by the **ospf** subsystem. The number of KB in use and unused are listed, followed by the percentage of the allocated memory that remains free. Generally, about 25% of the memory you allocated for the **ospf** subsystem (using the **`main getmem`** command) should be free. If the `Total Authorization Failures` field lists errors you need to allocate more memory. (See Section 2.5.2 on page 15.)

If the `Total Authorization Failures` field does not list errors, but the `Free` field indicates that less than 25% of the memory is free, you might need to allocate more memory. Periodically check the memory usage (by reissuing the **`stats`** command) to monitor the memory usage. If the percentage in the `Free` field continues to decrease, you need to allocate more memory.

3.5.2 Clearing OSPF Statistics

To clear OSPF statistics, issue the following command:

stats clear

Here is an example of this command.

```
23:PowerHub:ospf# stats clear
OSPF: Statistics Cleared.
```

The PowerHub software clears the counters for the statistics and begins collecting statistics again. Statistics also are cleared if OSPF routing is disabled, the software is rebooted, or the PowerHub system is powered down.

4 Route Exchange Commands

If you configure the PowerHub system as an Autonomous System Border router, the hub can exchange route information between RIP and OSPF. Networks that use RIP can access networks that use OSPF, and vice-versa.

This chapter describes the commands that let you define filters for exchanging routes between RIP and OSPF.

4.1 ACCESSING THE RIP SUBSYSTEM

To access the **rip** subsystem, issue the following command at the runtime prompt:

```
rip
```

4.2 RIP AND OSPF ROUTE EXCHANGE COMMANDS

Table 4–1 lists and describes the import filter commands in the **rip** subsystem and their syntax. For each command, the management capability is listed, as well as the section that contains additional information about the command.

TABLE 4–1 RIP import filter subsystem commands.

Command and Description	Capability*	See…
ospf-export-filter|oef add|a *<filnum>* **report|r|hide|h** *<addr> <mask>* Adds a filter for exporting routes from RIP to OSPF.	R	4.3.2
ospf-export-filter|oef chng|c *<filnum>* **report|r|hide|h [***<addr> <mask>***]** Changes an existing filter for exporting routes from RIP to OSPF.		
ospf-export-filter|oef del|d *<filter-list>* Deletes a filter for exporting routes from RIP to OSPF.		
ospf-export-filter|oef show|s *<filter-list>* Displays filters for exporting routes from RIP to OSPF.		
ospf-import-filter|oif add|a *<filnum>* **accept|a|discard|d** *<addr> <mask>* Adds a filter for importing routes from OSPF to RIP.	R	4.3.1
ospf-import-filter|oif chng|c *<filnum>* **accept|a|discard|d [***<addr> <mask>***]** Changes an existing filter for importing routes from OSPF to RIP.		
ospf-import-filter|oif del|d *<filter-list>* Deletes a filter for importing routes from OSPF to RIP.		
ospf-import-filter|oif show|s *<filter-list>* Displays filters for importing routes from OSPF to RIP.		

***R= Root.

4.3 FILTERING EXCHANGED ROUTES

The PowerHub software is configured to exchange all OSPF and RIP routes by default. If you do not want the PowerHub software to exchange all route information between RIP and OSPF, you can define filters to selectively block or forward routes exchanged between RIP and OSPF:

- To control the routes exported from RIP to OSPF, use the **rip ospf-export-filter** command. Define accept filters to explicitly export only the routes you specify. Define discard filters to export all routes except the routes you explicitly discard. (See Section 4.3.1 on page 65.)
- To control the routes imported from OSPF to RIP, use the **rip ospf-import-filter** command. Define accept filters to explicitly import only the routes you specify. Define discard filters to import all routes except the routes you explicitly discard. (See Section 4.3.2 on page 67.)

Routes are filtered according to the network address (or subnet address) and the subnet mask. Routes to the network or subnet are either accepted for import or export or discarded, based upon the action (accept or discard) you specify when you define the filter.

> **NOTE**: If both RIP and OSPF are enabled on the same interface, routing loops can occur. Accordingly, when you enable OSPF routing, the PowerHub software automatically turns off RIP listen on the IP interfaces that also are enabled as OSPF interfaces. You cannot enable RIP listen on IP interfaces that are enabled as OSPF interfaces.
>
> The PowerHub software does not change the setting for RIP talk.

4.3.1 Filtering the OSPF Routes Imported to RIP

The following sections describe the commands you use to add, display, change, or delete OSPF import filters.

4.3.1.1 Adding an OSPF Import Filter

To selectively block or forward the OSPF routes reported to RIP, use the **ospf-import-filter** command.

ospf-import-filter|oif add|a *<filnum>* **accept|a|discard|d** *<addr> <mask>*

where:

<filnum> Specifies the filter number. You can specify a number from **1** through **64**.

accept|a|discard|d

Specifies whether you want the PowerHub software to accept or discard routes to the networks corresponding to *<addr>* and *<mask>.*

If you specify **accept**, routes to only the networks specified by *<addr>* and *<mask>* are accepted. Routes to all other networks reported by OSPF are discarded. To accept additional routes, define additional filters to accept those routes.

Conversely, if you specify **discard**, routes to only the networks specified by *<addr>* and *<mask>* are discarded. Routes to all other networks reported by OSPF are accepted. To discard additional routes, define additional filters to discard those routes.

<addr> Specifies the IP network address or subnet address you want RIP to accept or discard. Use dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 through 9).

<mask> Specifies the IP subnet mask for the network or subnet you specify for *<addr>*. Use dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 through 9).

4.3.1.2 Displaying an OSPF Import Filter

To display an OSPF import filter, use the following command:

ospf-import-filter|oif show|s *<filter-list>*

where:

<filter-list> Specifies the filter number(s). You can specify a single filter number or a comma-separated list of filter numbers.

4.3.1.3 Changing an OSPF Import Filter

To change an existing OSPF import filter, use the following command:

ospf-import-filter|oif chng|c *<filnum>* **accept|a|discard|d**
[*<addr> <mask>***]**

The arguments are the same as those used by the **ospf-import-filter add** command. See Section 4.3.1.1 on page 65.

4.3.1.4 Deleting an OSPF Import Filter

To change an existing OSPF import filter, use the following command:

ospf-import-filter|oif del|d *<filter-list>*

where:

<filter-list> Specifies the filter number(s). You can specify a single filter number or a comma-separated list of filter numbers.

4.3.2 Filtering the RIP Routes Exported to OSPF

The following sections describe the commands you use to add, display, change, or delete OSPF import filters.

> **NOTE**: The PowerHub software exports all RIP routes into OSPF as type-2 external routes.

4.3.2.1 Adding an OSPF Export Filter

To selectively block or forward the RIP routes reported to OSPF, use the **ospf-export-filter** command.

ospf-export-filter|oif add|a *<filnum>* **accept|a|discard|d** *<addr> <mask>*

where:

<filnum> Specifies the filter number. You can specify a number from **1** through **64**.

accept|a|discard|d

Specifies whether you want the PowerHub software to accept or discard routes to the networks corresponding to *<addr>* and *<mask>.*

If you specify **accept**, routes to only the networks specified by *<addr>* and *<mask>* are accepted. Routes to all other networks reported by RIP are discarded. To accept additional routes, define additional filters to accept those routes.

Conversely, if you specify **discard**, routes to only the networks specified by *<addr>* and *<mask>* are discarded. Routes to all other networks reported by RIP are accepted. To discard additional routes, define additional filters to discard those routes.

<addr> Specifies the IP network address or subnet address you want OSPF to accept or discard. Use dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 through 9).

<mask> Specifies the IP subnet mask for the network or subnet you specify for *<addr>*. Use dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 through 9).

4.3.2.2 Displaying an OSPF Export Filter

To display an OSPF export filter, use the following command:

ospf-export-filter|oif show|s *<filter-list>*

where:

<filter-list> Specifies the filter number(s). You can specify a single filter number or a comma-separated list of filter numbers.

4.3.2.3 Changing an OSPF Export Filter

To change an existing OSPF export filter, use the following command:

ospf-import-filter|oif chng|c *<filnum>* **accept|a|discard|d [***<addr> <mask>***]**

The arguments are the same as those used by the **ospf-import-filter add** command. See Section 4.3.2.1 on page 67.

4.3.2.4 Deleting an OSPF Export Filter

To change an existing OSPF export filter, use the following command:

ospf-import-filter|oif del|d *<filter-list>*

where:

<filter-list>	Specifies the filter number(s). You can specify a single filter number or a comma-separated list of filter numbers.

Index

A

B

C

D

F

H

I

L

M

N

O

P

R

S

T

V